



7210 SAS X OS Quality of Service Guide

Software Version: 7210 SAS X OS 4.0 Rev. 01
October 2011

Document Part Number: 93-0383-01-01



This document is protected by copyright. Except as specifically permitted herein, no portion of the provided information can be reproduced in any form, or by any means, without prior written permission from Alcatel-Lucent.

Alcatel, Lucent, Alcatel-Lucent and the Alcatel-Lucent logo are trademarks of Alcatel-Lucent. All other trademarks are the property of their respective owners.

The information presented is subject to change without notice.

Alcatel-Lucent assumes no responsibility for inaccuracies contained herein.

Copyright 2011 Alcatel-Lucent. All rights reserved.

TABLE OF CONTENTS

Preface	11
Getting Started	
Alcatel-Lucent 7210 SAS-Series Services Configuration Process	16
QoS Policies	
QoS Overview	18
QoS Policies	19
Service and Network QoS Policies	22
Network QoS Policies	23
Network Queue QoS Policies	28
Meter Parameters	29
Queue Parameters	35
Service Ingress QoS Policies	41
Hierarchical Ingress Policing	45
Service Egress QoS Policies	46
Access Egress QoS Policies	51
Buffer Pools	52
Queue Management Policies	53
Queue Management Policy Parameters	59
Remark Policy	60
Egress Port Rate Limiting	61
Forwarding Classes	62
QoS Policy Entities	64
Configuration Notes	65
Port Level Egress Rate-Limiting	
Overview	68
Applications	68
Affect of Port Level Rate-Limiting on Queue Functionality	69
Basic Configurations	70
Modifying Port Level Egress-Rate Command	71
Removing Port Level Egress-Rate Command	72
Default Egress-Rate Values	72
Port Level Egress-Rate Command Reference	73
Frame Based Accounting	
Overview	78
Affects of Enabling Ingress Frame Based Accounting on Ingress Meter Functionality	78
Affects of Egress Frame Based Accounting on Queue Functionality	78
Accounting and Statistics	78
Basic Configurations	79
Enabling and Disabling Frame-Based Accounting	80
Default Frame-Based-Accounting Values	80
Frame Based Accounting Command Reference	81
Configuration Commands	82

Table of Contents

Show Commands	83
Network QoS Policies	
Overview	88
Normal QoS Operation	90
Network Qos Policy (ip-interface type) Functionality	91
Upgrading from Release 3.0 to Release 4.0	93
DSCP Marking CPU Generated Traffic	94
Default DSCP Mapping Table	95
Basic Configurations	96
Create a Network QoS Policy	96
Default Network Policy Values	99
Service Management Tasks	103
Deleting QoS Policies	103
Remove a Policy from the QoS Configuration	104
Copying and Overwriting Network Policies	104
Editing QoS Policies	105
Resource Allocation for Network QoS policy	106
Network QoS Policies Resource Usage Examples	108
Network QoS Policy Command Reference	115
Network Queue QoS Policies	
Overview	150
Basic Configurations	152
Create a Network Queue QoS Policy	152
Applying Network Queue Policies	154
Ethernet Ports	154
Default Network Queue Policy Values	155
Service Management Tasks	158
Deleting QoS Policies	158
Copying and Overwriting QoS Policies	159
Editing QoS Policies	161
Network Queue QoS Policy Command Reference	163
Service Ingress QoS Policies	
Overview	176
Default SAP Ingress Policy	177
SAP Ingress Policy Defaults	178
Service Ingress Meter Selection Rules	179
Service Ingress Policy Configuration Considerations	180
Basic Configurations	186
Create Service Ingress QoS Policies	186
Service Ingress QoS Policy	187
Applying Service Ingress Policies	220
Service Management Tasks	222
Deleting QoS Policies	222
Remove a QoS Policy from Service SAP(s)	222
Copying and Overwriting QoS Policies	223
Remove a Policy from the QoS Configuration	224

Editing QoS Policies	224
Service SAP QoS Policy Command Reference	225
Access Egress QoS Policies	
Overview	262
Basic Configurations	262
Create Access Egress QoS Policies	263
Access Egress QoS Policy	263
Applying Access Egress QoS Policies	264
Default Access Egress QoS Policy Values	265
Editing QoS Policies	266
Deleting QoS Policies	267
Removing a Policy from the QoS Configuration	267
Access Egress QoS Policy Command Reference	269
SAP Egress Policies	
Overview	278
Configuration Guidelines	279
Basic Configurations	279
Create a SAP Egress Policy	279
Editing QoS Policies	282
SAP Egress Policy Command Reference	297
QoS Schedulers	
Overview	300
Queue Management Policies	
Overview	304
Basic Configurations	305
Creating a Queue Management Policy	305
Editing QoS Policies	306
Queue Management Policy Command Reference	307
Remark Policies	
Overview	320
Basic Configurations	322
Creating a Remark Policy	322
Editing QoS Policies	323
Remark Policy Command Reference	325
Multipoint Bandwidth Management	
Overview	344
Configuration Guidelines	345
Multipoint Bandwidth Management Command Reference	347
Standards and Protocol Support	361
Index	365

Table of Contents

LIST OF TABLES

Getting Started

Table 1:	Configuration Process	16
----------	---------------------------------	----

QoS Policies

Table 2:	QoS Policy Types and Descriptions	21
Table 3:	QoS Policy Types and Descriptions	21
Table 4:	Default Network QoS Policy (type = ip-interface) Egress Marking	24
Table 5:	Default Network QoS Policy (type = ip-interface) to FC Mapping	25
Table 6:	Default Network QoS Policy of type 'port' Egress Marking	26
Table 7:	Default Network QoS Policy of Type Port - Dot1p/DSCP to FC Mapping	27
Table 8:	Default Network Queue Policy Definition.	28
Table 9:	Supported Hardware rates and burst step sizes for CIR and PIR values	31
Table 10:	Administrative Rate Example	32
Table 11:	Supported Hardware Rates and CIR/PIR Values	38
Table 12:	Service Ingress QoS Policy IP Match Criteria	43
Table 13:	Service Ingress QoS Policy MAC Match Criteria	43
Table 14:	MAC Match Ethernet Frame Types.	43
Table 15:	MAC Match Criteria Frame Type Dependencies.	44
Table 16:	Default Service Ingress Policy ID 1 Definition.	44
Table 17:	Default SAP Egress Policy ID 1 Definition	48
Table 18:	TAF Impact on Shared Buffer Average Utilization Calculation	57
Table 19:	Default Slope Policy Definition	59
Table 20:	Slope Policy Defaults	60
Table 21:	Forwarding Classes	62
Table 22:	Forwarding Class to Queue-ID Map	63

Network QoS Policies

Table 23:	DSCP and Dot1p Marking.	94
Table 24:	Network Policy Defaults	99
Table 25:	Default Network QoS Policy of Type IP Interface, LSP EXP to FC Mapping on Ingress (Color aware policing is supported on network ingress.)100	
Table 26:	Default DSCP Names to DSCP Value Mapping Table	131
Table 27:	Default Class Selector Code Points to DSCP Value Mapping Table	131
Table 28:	Show QoS Network Output Fields	143
Table 29:	Show QoS Network Output Fields	147

Network Queue QoS Policies

Table 30:	pir-level Assignment to queue based on the cir-level	151
Table 31:	Network Queue Labels and Descriptions	172

Service Ingress QoS Policies

Table 32:	SAP Ingress Policy Defaults	178
-----------	---------------------------------------	-----

SAP Egress Policies

Table 33:	FC to Queue Map	278
-----------	---------------------------	-----

QoS Schedulers

Table 34:	pir-level Assignments	300
-----------	---------------------------------	-----

Queue Management Policies

Table 35:	Show Queue Management Policy Output Fields	315
-----------	--	-----

Remark Policies

Table 36:	Summary of remark policy and attachment points	320
Table 37:	Show Remark Policy Output Fields	334

Multipoint Bandwidth Management

Table 38:	FC Queue Table	351
Table 39:	Show Multipoint-management Policy Output Fields	354

LIST OF FIGURES

QoS Policies

Figure 1:	7210 SAS X Traffic Types	22
Figure 2:	Traffic Queuing Model for Forwarding Classes	42
Figure 3:	RED Slope Characteristics	56

Preface

About This Guide

This guide describes the Quality of Service (QoS) provided by the 7210-SAS X OS and presents examples to configure and implement various protocols and services.

This document is organized into functional chapters and provides concepts and descriptions of the implementation flow, as well as Command Line Interface (CLI) syntax and command usage.

Audience

This manual is intended for network administrators who are responsible for configuring the 7210 SAS-Series routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Protocols, standards, and services described in this manual include the following:

- CLI concepts
- Quality of Service (QoS) policies and profiles

List of Technical Publications

The 7210-SAS M, X OS documentation set is composed of the following books:

- 7210-SAS M, X OS Basic System Configuration Guide
This guide describes basic system configurations and operations.
- 7210-SAS M, X OS System Management Guide
This guide describes system security and access configurations as well as event logging and accounting logs.
- 7210-SAS M, X OS Interface Configuration Guide
This guide describes card, Media Dependent Adapter (MDA), and port provisioning.
- 7210-SAS M, X OS Router Configuration Guide
This guide describes logical IP routing interfaces and associated attributes such as an IP address, port, link aggregation group (LAG) as well as IP and MAC-based filtering.
- 7210 SAS X OS Services Guide
This guide describes how to configure service parameters such as customer information, and user services.
- 7210-SAS M, X OS OAM and Diagnostic Guide
This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.
- 7210 SAS X OS Quality of Service Guide
This guide describes how to configure Quality of Service (QoS) policy management.
- 7210-SAS M, X OS MPLS Guide
This guide describes how to configure Multiprotocol Label Switching (MPLS) and Label Distribution Protocol (LDP).
- 7210-SAS M, X OS OS Routing Protocols Guide
This guide provides an overview of routing concepts and provides configuration examples for OSPF, IS-IS, and route policies.

Technical Support

If you purchased a service agreement for your 7210 SAS device and related products from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance. If you purchased an Alcatel-Lucent service agreement, contact your welcome center.

Web: http://www1.alcatel-lucent.com/comps/pages/carrier_support.jhtml

Getting Started

In This Chapter

This chapter provides process flow information to configure Quality of Service (QoS) policies and provision services.

Alcatel-Lucent 7210 SAS-Series Services Configuration Process

[Table 1](#) lists the tasks necessary to configure and apply QoS policies. This guide is presented in an overall logical configuration flow. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

Table 1: Configuration Process

Area	Task	Chapter
Policy configuration	Configuring QoS Policies	
	• Egress Rate	Port Level Egress Rate-Limiting on page 67
	• Accounting Mode	Frame Based Accounting on page 77
	• Network	Network QoS Policies on page 87
	• Network queue	Network Queue QoS Policies on page 149
	• SAP ingress	Service Ingress QoS Policies on page 175
	• SAP egress	SAP Egress Policies on page 277
	• Scheduler Policies	QoS Schedulers on page 299
	• Queue Management	Queue Management Policies on page 303
	• Remark Policies	Remark Policies on page 319
Reference	• Multipoint Bandwidth Management	Multipoint Bandwidth Management on page 343
	• List of IEEE, IETF, and other proprietary entities	Standards and Protocol Support on page 361

QoS Policies

In This Chapter

This chapter provides information about Quality of Service (QoS) policy management.

Topics in this chapter include:

- [QoS Overview on page 18](#)
- [Service and Network QoS Policies on page 22](#)
 - [Network QoS Policies on page 23](#)
 - [Network Queue QoS Policies on page 28](#)
 - [Service Ingress QoS Policies on page 41](#)
 - [Queue Parameters on page 35](#)
- [Queue Management Policies on page 53](#)
- [QoS Policy Entities on page 64](#)
- [Configuration Notes on page 65](#)

QoS Overview

The 7210 SAS X is designed with QoS mechanisms on both ingress and egress to support multiple services per physical port. The 7210 SAS X has extensive and flexible capabilities to classify, police, shape, and mark traffic.

In the Alcatel-Lucent service router's service model, a service is provisioned on the provider-edge (PE) equipment. Service data is encapsulated and then sent in a service tunnel to the far-end Alcatel-Lucent service router where the service data is delivered.

The operational theory of a service tunnel is that the encapsulation of the data between the two Alcatel Lucent service routers appear like a Layer 2 path to the service data although it is really traversing an IP or IP/MPLS core. The tunnel from one edge device to the other edge device is provisioned with an encapsulation and the services are mapped to the tunnel that most appropriately supports the service needs.

The 7210 SAS supports eight forwarding classes internally named: Network-Control, High-1, Expedited, High-2, Low-1, Assured, Low-2 and Best-Effort. The forwarding classes are discussed in more detail in [Forwarding Classes on page 62](#).

7210 SAS use QoS policies to control how QoS is handled at distinct points in the service delivery model within the device. There are different types of QoS policies that cater to the different QoS needs at each point in the service delivery model. QoS policies are defined in a global context in the 7210 SAS and only take effect when the policy is applied to a relevant entity.

QoS policies are uniquely identified with a policy ID number or name. Policy ID 1 or Policy ID "default" is reserved for the default policy which is used if no policy is explicitly applied.

The QoS policies within the 7210 SAS can be divided into three main types:

- QoS policies are used for classification, ingress policing, egress queue attributes, and marking.
- Queue management policies define buffer allocations and WRED slope definitions.
- Scheduler policies determine how queues are scheduled.

QoS Policies

7210 SAS X QoS policies are applied on service ingress, service egress, network port ingress and egress, and network IP interfaces.

- Classification rules for how traffic is mapped to forwarding classes
- Forwarding class association with meters and meter parameters used for policing (rate-limiting).
- Queuing parameters for shaping, scheduling, and buffer allocation
- QoS marking/interpretation

There are several types of QoS policies:

- Service ingress
- Service egress
- Network (for ingress and egress)
- Network queue (for egress)
- Scheduler
- Queue Management
- Remark policies

Service ingress QoS policies are applied to the customer-facing Service Access Points (SAPs). Traffic that enters through the SAP is classified to map it to a Forwarding Class (FC). Forwarding class is associated with meters on ingress. The mapping of traffic to meters can be based on combinations of customer QoS marking (IEEE 802.1p bits), IP and MAC criteria. The characteristics of the forwarding class meters are defined within the policy as to the number of forwarding class meters for unicast traffic and the meter characteristics (like CIR, PIR, etc.). Each of the forwarding classes can be associated with different unicast parameters. A service ingress QoS policy also defines up to three (3) meters per forwarding class to be used for multipoint traffic for multipoint services. There can be up to 16 meters in total per Service ingress QoS policies. In the case of the VPLS, four types of forwarding are supported (which is not to be confused with forwarding classes); unicast, multicast, broadcast, and unknown. Multicast, broadcast, and unknown types are flooded to all destinations within the service while the unicast forwarding type is handled in a point-to-point fashion within the service.

Service egress QoS policies are applied to SAPs and map forwarding classes to service egress queues for a service. The system allocates 8 queues per SAP for the 8 forwarding classes. A service egress QoS policy also defines how to remark the forwarding class to IEEE 802.1p bits in the customer traffic.

There are two types of network QoS policies, one applied to a network IP interface and the other type is applied to a network port. On ingress, the policy applied to an IP interface maps incoming

values to forwarding class and profile state for the traffic received from the core network. On egress, the policy maps forwarding class and profile state to values for traffic to be transmitted into the core network. The network policy applied to a network port maps incoming IP packets, DSCP or Dot1p values, to the forwarding class and the profile state for the traffic received from the core network. On egress, the policy maps forwarding class and profile state to DSCP and/or Dot1p values for IP traffic to be transmitted into the core network.

Network queue policies are applied on egress to ports. The policies define the forwarding class queue characteristics for these entities. The FCs are mapped onto the queues. There are 8 queues at the port level. FC-to-queue mapping is static and is not configurable. The number of queues are static and service are always 8 queues at the port level.

Service ingress, service egress, and network QoS policies are defined with a scope of either *template* or *exclusive*. Template policies can be applied to multiple entities (such as SAPs and ports) whereas exclusive policies can only be applied to a single entity.

One service ingress and one service egress QoS policy can be applied to a specific SAP. One network QoS policy can be applied to a specific IP interface or network port based on the type of network QoS policy. A network QoS policy defines both ingress and egress behavior. One network queue policy can be applied to the network port.

If no QoS policy is explicitly applied to a SAP, port or interface, a default QoS policy is applied.

A summary of the major functions performed by the QoS policies is listed in [Table 2](#).

Table 2: QoS Policy Types and Descriptions

Policy Type	Applied at...	Description	Page
Service Ingress	SAP ingress	• Defines up to 32 forwarding class meters and meter parameters for traffic classification. • Defines match criteria to map flows to the meters based on any one of the criteria (IP or MAC).	41
Service Egress	SAP Egress	• Defines up to 8 forwarding class queues. • Maps forwarding classes to the queues. • Define Queue parameters for the queues. • Defines FC to remarking values. • Defines CIR levels and PIR weights that determines how the queue gets prioritized by the scheduler.	41
Network	IP interface	Packets are marked using QoS policies on edge devices. Invoking a QoS policy on a network port allows for the packets that match the policy criteria to be remarked. • Used for classification/marketing of MPLS packets. • At ingress, defines MPLS LSP-EXP to FC mapping and 12 meters used by FCs. • At egress, defines FC to MPLS LSP-EXP marking.	23
Network	Ports	• Used for classification/marketing of IP packets. • At ingress, defines DSCP or Dot1p to FC mapping and 8 meters. • At egress, defines FC to DSCP or Dot1p marking or both.	
Network Queue	Network ports	• Defines forwarding class mappings to network queues and queue characteristics for the queues.	28
Queue Management Policies	Queues at service egress and network egress	• Defines the CBS and MBS parameters for the queues. • Enables or disables the high-slope and low-slope parameters for the queues.	53
Remark	SAP egress Network egress	• Defines the forwarding class to remarking values.	94

Service and Network QoS Policies

The QoS mechanisms within the 7210 SAS X are specialized for the type of traffic on the interface. For customer interfaces, there is service ingress and egress traffic, and for IP interfaces, there is network ingress and network egress traffic ([Figure 1](#)).




Figure 1: 7210 SAS X Traffic Types

The 7210 SAS uses QoS policies applied to a SAP for a service or to an network IP interface or a network port to define the queuing, queue attributes, and QoS marking/interpretation.

The 7210 SAS supports four types of service and network QoS policies:

- Service ingress QoS policies
- Service egress QoS policies
- Network QoS policies
- Network Queue QoS policies

Network QoS Policies

Two types of network QoS policies can be defined, **ip-interface** and **port**. By default, when a network QoS policy is created, it is an **ip-interface** type.

A network QoS policy of type **ip-interface** is created in the **configure>qos>network network-policy-id create** context.

A network QoS policy of type **port** is created in the **configure>qos>network network-policy-id network-policy-type port create** context.

When a network QoS policy of type **ip-interface** is applied to IP interface, it is used for classification of MPLS packets based on LSP-EXP bits.

When a network QoS policy of type **port** is applied to port, it is used for classification of IP packets based on DSCP or Dot1p bits.

Network QoS policies (**ip-interface** type) define ingress forwarding class meters and maps traffic to those meters for IP interfaces. When a network QoS policy is created, it always has two meters defined that cannot be deleted, one for the all unicast traffic and one for all multipoint traffic. These meters exist within the definition of the policy. The meters only get instantiated in hardware when the policy is applied to an IP interface. A remarking policy can be specified to define the forwarding class to EXP bit marking, on the egress.

- Ingress
 - Defines EXP value mapping to forwarding classes.
 - Defines forwarding class to meter mapping.
- Egress
 - Specifies a remark policy that defines the forwarding class to EXP value markings.
 - Remarking of QoS bits can be enabled or disabled.

The required elements to be defined in a network QoS policy are:

- A unique network QoS policy ID.
- Specifies a remark policy to define the forwarding class to value mappings for each forwarding class.
- A default ingress forwarding class and in-profile/out-of-profile state.
- At least one default unicast forwarding class meter. The parameters that can be configured for a meter are discussed in [Meter Parameters on page 29](#).
- At least one multipoint forwarding class meter.

Optional network QoS policy elements include:

- Additional unicast meters up to the maximum number allowed for network ingress.
- Additional multipoint meters up to the maximum number allowed for network ingress.
- EXP value to forwarding class and profile state mappings for all EXP values received.

Network policy ID 2 is reserved as the default network QoS policy of type IP interface. The default policy cannot be deleted or changed.

Default network QoS policy 2 is applied to all IP interfaces which do not have another network QoS policy explicitly assigned.

Note that FC to Dot1p marking is used to mark IP and MPLS traffic sent out through that port, if marking is enabled and remark policy specifies the values for both.

The network QoS policy applied at network egress (for example, on an IP interface) determines how or if the profile state is marked in packets transmitted into the service core network. If the profile state is marked in the service core packets, out-of-profile packets are preferentially dropped over in-profile packets at congestion points in the core network. For network egress, traffic remarking in the network QoS policy is disabled. [Table 5](#) lists the default mapping of forwarding class to EXP values.

Table 4: Default Network QoS Policy (type = ip-interface) Egress Marking

FC-ID	FC Name	FC Label	Egress EXP Marking	
			In-Profile	Out-of-Profile
7	Network Control	nc	111 - 7	111 - 7
6	High-1	h1	110 - 6	110 - 6
5	Expedited	ef	101 - 5	101 - 5
4	High-2	h2	100 - 4	100 - 4
3	Low-1	l1	011 - 3	010-2
2	Assured	af	011-3	010 - 2
1	Low-2	l2	001 - 1	001 - 1
0	Best Effort	be	000 - 0	000 - 0

For network ingress, [Table 5](#) lists the default mapping of EXP values to forwarding class and profile state for the default network QoS policy.

Table 5: Default Network QoS Policy (type = ip-interface) to FC Mapping

Value	7210 FC Ingress	Profile
0	be	Out
1	l2	In
2	af	Out
3	af	In
4	h2	In
5	ef	In
6	h1	In
7	nc	In

“port” Type Network QoS Policy

Network QoS policy of type **port** defines ingress forwarding class meters and maps traffic to those meters for only IP traffic received on network ports. When a network policy of this type is created it has a single unicast meter which cannot be deleted. These meters exist within the definition of the policy. The meters get instantiated in hardware, only when the policy is applied to a network port. It also defines the forwarding class to DSCP and/or Dot1p marking to be used for packets sent out through that port.

A network QoS policy of type port defines both the ingress and egress handling of QoS on the network port.

The following functions are defined:

- Ingress
 - Defines DSCP or Dot1p value mapping to forwarding classes. Only one type supported, such as DSCP or Dot1p, per policy.
 - Defines forwarding class to meter mapping.
- Egress
 - Specifies remark policy that defines forwarding class to DSCP or Dot1p (or both) value markings.
 - Remark of QoS bits is always disabled

The required elements to be defined in a network QoS policy of port type are:

- A unique network QoS policy ID and network-policy-type set to **port**.
- Egress forwarding class to DSCP or Dot1p (or both) value mappings for each forwarding class.
- A default ingress forwarding class and in-profile/out-of-profile state.
- At least one default unicast forwarding class meter. The parameters that can be configured for a meter are discussed in Meter Parameters on page 25.

Optional network QoS policy elements include:

- Additional unicast meters up to a total of 8.
- A DSCP or Dot1p (or both) value to forwarding class and profile state mappings for all DSCP or Dot1p values received.

Network policy ID 1 is reserved as the default network QoS policy of type port. The default policy cannot be deleted or changed.

The default network QoS policy is applied to all network ports which do not have another network QoS policy explicitly assigned.

[Table 6](#) lists the default mapping of forwarding class to Dot1p and DSCP values.

Table 6: Default Network QoS Policy of type 'port' Egress Marking

FC-ID	FC Name	FC Label	Egress DSCP Marking		Egress Dot1p Marking	
			In-Profile	Out-of-Profile	In-Profile	Out-of-profile
7	Network Control	nc	nc2	nc2	111 - 7	111 - 7
6	High-1	h1	nc1	nc1	110-6	110-6
5	Expedited	ef	ef	ef	101-5	101-5
4	High-2	h2	af41	af41	100-4	100-4
3	Low-1	l1	af21	af22	011-3	010-2
2	Assured	af	af11	af12	011-3	010-2
1	Low-2	l2	cs1	cs1	001-1	001-1
0	Best Effort	be	be	be	000-0	000-0

For network ingress, [Table 7](#) lists the default mapping of Dot1p/DSCP values to forwarding class and profile state for the default network QoS policy of type port. Color aware policing is supported on network ingress.

Table 7: Default Network QoS Policy of Type Port - Dot1p/DSCP to FC Mapping

DSCP Value	Dot1p Value	FC Ingress	Profile
be	0	be	In, Out
cs1	1	l2	In, Out
af12	2	af	Out
af11	3	af	In
af41	4	h2	In, Out
ef	5	ef	In, Out
nc1	6	h1	In, Out
nc2	7	nc	In, Out

Network Queue QoS Policies

Network queue policies define the network forwarding class queue characteristics. Network queue policies are applied on egress on network ports. The system allocates 8 queues for the network port and FCs are mapped to these 8 queues. FC to queue mapping is not a configurable entity. All policies will use eight queues like the default network queue policy.

The queue characteristics that can be configured on a per-forwarding class basis are:

- Peak Information Rate (PIR) as a percentage of egress port bandwidth
- Committed Information Rate (CIR) as a percentage of egress port bandwidth
- Committed burst size (CBS)
- Maximum burst size (MBS)
- CIR-Level and PIR-Weight
- Adaptation rules for CIR/PIR
- WRED Slope Parameters (using the queue-management policy)

Network queue policies are identified with a unique policy name which conforms to the standard 7210 SAS alphanumeric naming conventions.

The system default network queue policy is named **default** and cannot be edited or deleted. CBS values cannot be provisioned. CBS is set to 750KB and MBS is set to 1750KB for all the queues. [Table 8](#) describes the default network queue policy definition.

Table 8: Default Network Queue Policy Definition.

Forwarding Class	Queue	Definition
Network-Control (nc)	Queue 8	• PIR = 10% • CIR = 10% • CIR-Level 8 • PIR-Weight 1
High-1 (h1)	Queue 7	• PIR = 100% • CIR = 10% • CIR-Level 7 • PIR-Weight 1
Expedited (ef)	Queue 6	• PIR = 100% • CIR = 100% • CIR-Level 6 • PIR-Weight 1

Table 8: Default Network Queue Policy Definition. (Continued)

Forwarding Class	Queue	Definition (Continued)
High-2 (h2)	Queue 5	• PIR = 100% • CIR = 100% • CIR-Level 5 • PIR-Weight 1
Low-1 (l1)	Queue 4	• PIR = 100% • CIR = 25% • CIR-Level 4 • PIR-Weight 1
Assured (af)	Queue 3	• PIR = 100% • CIR = 25% • CIR-Level 3 • PIR-Weight 1
Low-2 (l2)	Queue 2	• PIR = 100% • CIR = 25% • CIR-Level 2 • PIR-Weight 1
Best-Effort (be)	Queue 1	• PIR = 100% • CIR = 0% • CIR-Level 1 • PIR-Weight 1

Meter Parameters

This section describes the meter parameters provisioned on access and network meters provisioned on IP interfaces for QoS.

The meter parameters are:

- [Meter ID on page 30](#)
- [Committed Information Rate on page 30](#)
- [Peak Information Rate on page 30](#)

- [Adaptation Rule for Meters on page 31](#)
 - [Committed Burst Size on page 32](#)
 - [Maximum Burst Size on page 32](#)
 - [Meter Counters on page 33](#)
 - [Meter Modes on page 33](#)
-

Meter ID

The meter ID is used to uniquely identify the meter. The meter ID is only unique within the context of the QoS policy within which the meter is defined.

Committed Information Rate

The committed information rate (CIR) for a meter is the long term average rate at which traffic is considered as conforming traffic or in-profile traffic. The higher the rate, the greater the throughput user can expect. The user will be able to burst above the CIR and up to PIR for brief periods of time. The time and profile of the packet is decided based on the burst sizes as explained in the following sections.

When defining the CIR for a meter, the value specified is the administrative CIR for the meter. The 7210 SAS X has a number of native rates in hardware that it uses to determine the operational CIR for the meter. The user has some control over how the administrative CIR is converted to an operational CIR should the hardware not support the exact CIR and PIR combination specified. Refer to the interpretation of the administrative CIR in [Adaptation Rule for Meters on page 31](#).

The CIR for meter is provisioned on service ingress and network ingress within service ingress QoS policies and network QoS policies, respectively.

Peak Information Rate

The peak information rate (PIR) defines the maximum rate at which packets are allowed to exit the meter. It does not specify the maximum rate at which packets may enter the meter; this is governed by the meter's ability to absorb bursts and is defined by its maximum burst size (MBS).

When defining the PIR for a meter, the value specified is the administrative PIR for the meter. The 7210 SAS X has a number of native rates in hardware that it uses to determine the operational PIR for the meter. The user has some control over how the administrative PIR is converted to an operational PIR should the hardware not support the exact CIR and PIR combination specified. Refer to the interpretation of the administrative PIR in [Adaptation Rule for Meters on page 31](#).

The PIR for meter is provisioned on service ingress and network ingress within service ingress QoS policies and network QoS policies, respectively

Adaptation Rule for Meters

The adaptation rule provides the QoS provisioning system with the ability to adapt the administrative rates provisioned for CIR and PIR, to derive the operational rates based on the underlying capabilities of the hardware. The administrative CIR and PIR rates are translated to actual operational rates enforced by the hardware meter. The rule provides a constraint, when the exact rate is not available due to hardware capabilities.

The hardware rate step-size is provided in table [Table 9](#):

Table 9: Supported Hardware rates and burst step sizes for CIR and PIR values

Rate (kbits_sec)	Burst (kbits_burst)	Rate Step Size (bits)	Burst Step Size (bits)
0-4194296	0-16773	8000	4096
4194297-8388592	16774-33546	16000	8192
8388593-16777184	33547-67092	32000	16384
16777185-33554368	67093-134184	64000	32768
33554369-67108736	134185-268369	128000	65536
67108737-134217472	268370-536739	256000	131072
134217473-268434944	536739-1073479	512000	262144
268434945-536869888	1073480-2146959	1024000	524288

The system attempts to find the best operational rate depending on the defined constraint. The supported constraints are listed below:

- Minimum: Find the next multiple of step-size that is equal to or greater than the specified rate.
- Maximum: Find the next multiple of step-size that is equal to or less than the specified rate.
- Closest: Find the next multiple of step-size that is closest to the specified rate.

[Table 10](#) lists the rate values configured in the hardware when different PIR or CIR rates are specified in the CLI.

Table 10: Administrative Rate Example

Administrative Rate	Operation Rate (Min)	Operation Rate (Max)	Operation Rate (Closest)
8	8	8	8
10	16	8	8
118085	11808	11800	11808
46375	46376	46368	46376

If user has configured any value greater than 0 and less than 8 then operation rate configured on hardware would be 8 kbps irrespective of the constraint used.

Note: The burst size configured by the user affects the rate step-size used by the system. The system uses the step size in a manner that both the burst-size and rate parameter constraints are met. For example, if the rate specified is less than 4Gbps but the burst size configured is 17Mbits, then the system uses rate step-size of 16Kbits and burst step-size of 8192bits (refer to [Table 9](#), row#2)

Committed Burst Size

The committed burst size parameter specifies the maximum burst size that can be transmitted by the source at the CIR while still complying with the CIR. If the transmitted burst is lower than the CBS value then the packets are marked as in-profile by the meter to indicate that the traffic is complying meter configured parameters.

The operational CBS set by the system is adapted from the user configured value by using the minimum constraint.

Maximum Burst Size

For trTCM, the maximum burst size parameter specifies the maximum burst size that can be transmitted by the source at the PIR while complying with the PIR. If the transmitted burst is lower than the MBS value then the packets are marked as out-profile by the meter to indicate that the traffic is not complying with CIR, but complying with PIR.

For srTCM, the maximum burst size parameter specifies the maximum burst size that can be transmitted by the source while not complying with the CIR. If the transmitted burst is lower than the MBS value then the packets are marked as out-profile by the meter to indicate that the traffic is not complying with CIR.

If the packet burst is higher than MBS then packets are marked as red and dropped.

The operational MBS set by the system is adapted from the user configured value by using the minimum constraint.

Meter Counters

The 7210 SAS X maintains following counters for meters within the system for granular billing and accounting. Each meter maintains the following counters:

- Counters for packets or octets marked as in-profile by meter
 - Counters for packets or octets marked as out-of-profile by meter
-

Meter Modes

The 7210 SAS X supports following meter modes:

- srtcm: Single Rate Three Color Marking
- trtcm: Two Rate Three Color Marking
- trtcm1: Two Rate Three Color Marking1 (Applicable only for Service Ingress QoS Policies)
- trtcm2: Two Rate Three Color Marking2 (Applicable only for Service Ingress QoS Policies)

In srtcm the CBS and MBS Token buckets are replenished at single rate, that is, CIR where as in case of trtcm CBS and MBS buckets are individually replenished at CIR and PIR rates, respectively. trtcm1 implements the policing algorithm defined in RFC2698 and trtcm2 implements the policing algorithm defined in RFC4115.

Color Aware Policing

The 7210 SAS X supports Color Aware policing at the network ingress, where as at service ingress policing is color blind. In color aware policing user can define the color of the packet using the classification and feed those colored packets to the meter. A color aware meter would treat those packets with respect to the color defined.

- If the packet is pre-colored as in-profile (or also called as Green colored packets) then depending on the burst size of the packet meter can either mark it in-profile or out-profile.

- If the packet is pre-colored as out-profile (also called as Yellow colored packets) then even if the packet burst is lesser than the current available CBS, it would not be marked as in-profile and remain as out-profile.
- If the packet burst is higher than the MBS then it would be marked as Red and would be dropped by meter at ingress.

The profile marked by the meter is used to determine the packets eligibility to be enqueued into a buffer at the egress (when a slope policy is configured at the egress).

Queue Parameters

This section describes the queue parameters provisioned on service queues and network queues for QoS.

The queue parameters are:

- [Queue ID on page 35](#)
 - [Committed Information Rate on page 36](#)
 - [Peak Information Rate on page 37](#)
 - [Adaptation Rule for Queues on page 38](#)
 - [Committed Burst Size and the Maximum Burst Size \(MBS\) on page 40](#)
-

Queue ID

The queue ID is used to uniquely identify the queue. The queue ID is only unique within the context of the QoS policy within which the queue is defined. On 7210 SAS X, the queue ID is not a user configurable entity but the queue ID is statically assigned to the 8 Queues on the port according to FC-QID map table shown in [Table 22](#).

Committed Information Rate

The committed information rate (CIR) for a queue performs two distinct functions:

1. Minimum bandwidth guarantees — ueues CIR setting provides the bandwidth which will be given to this queue as compared to other queues on the port competing for a share of the available link bandwidth. The queue CIR does not necessarily guarantee bandwidth in all scenarios and also depends on factors such as CIR oversubscription and link port bandwidth capacity. For each packet in an egress queue, the CIR is checked with the current transmission rate of the queue. If the current rate is at or below the CIR threshold, the queue is considered in-profile. If the current rate is above the threshold, the queue is considered out-of-profile. This in and out profile state of queue is linked to scheduler prioritizing behavior as discussed below.
2. Scheduler queue priority metric — The scheduler serving a group of egress queues prioritizes individual queues based on their current CIR and PIR states. Queues operating below their CIR are always served before those queues operating at or above their CIR.

Queues at the egress never marks the packets as in-profile or out-profile based on the queue CIR, PIR values. The in-profile and out-profile state of the queue interacts with the scheduler mechanism and provides the minimum and maximum bandwidth guarantees.

When defining the CIR for a queue, the value specified is the administrative CIR for the queue. The user has some control over how the administrative CIR is converted to an operational CIR should the hardware not support the exact CIR and PIR combination specified. The interpretation of the administrative CIR is discussed below in [Adaptation Rulefor Queues on page 38](#)

Although the 7210 SAS is flexible in how the CIR can be configured, there are conventional ranges for the CIR based on the forwarding class of a queue. A queue associated with the high-priority class normally has the CIR threshold equal to the PIR rate although the 7210 SAS allows the CIR to be provisioned to any rate below the PIR should this behavior be required.

The CIR for a queue is provisioned on egress within service egress QOS policy.

The CIR for the network queues are defined within network queue policies based on the forwarding class. The CIR for the network queues is specified as a percentage of the network interface bandwidth.

Peak Information Rate

The peak information rate (PIR) defines the maximum rate at which packets are allowed to exit the queue. It does not specify the maximum rate at which packets may enter the queue; this is governed by the queue's ability to absorb bursts. The actual transmission rate of an egress queue depends on more than just its PIR. Each queue is competing for transmission bandwidth with other queues. Each queue's PIR, CIR and the relative priority and/or weight of the scheduler serving the queue, all combine to affect a queue's ability to transmit packets.

The PIR is provisioned on service egress queues within service egress QoS policies.

The PIR for network queues are defined within network queue policies based on the forwarding class. The PIR for the queues is specified as a percentage of the network interface bandwidth.

When defining the PIR for a queue, the value specified is the administrative PIR for the queue. The user has some control over how the administrative PIR is converted to an operational PIR should the hardware not support the exact CIR and PIR values specified. The interpretation of the administrative PIR is discussed below in [Adaptation Rule for Queues on page 38](#)

Adaptation Rulefor Queues

The adaptation rule provides the QoS provisioning system with the ability to adapt specific CIR and PIR defined administrative rates to the underlying capabilities of the hardware the queue will be created on to derive the operational rates. The administrative CIR and PIR rates are translated to actual operational rates enforced by the hardware queue. The rule provides a constraint used when the exact rate is not available due to hardware implementation trade-offs.

For the CIR and PIR parameters individually, the system will attempt to find the best operational rate depending on the defined constraint. The supported constraints are:

- Minimum — Find the hardware supported rate that is equal to or higher than the specified rate.
- Maximum — Find the hardware supported rate that is equal to or lesser than the specified rate.
- Closest — Find the hardware supported rate that is closest to the specified rate.

Depending on the hardware upon which the queue is provisioned, the actual operational CIR and PIR settings used by the queue will be dependant on the method the hardware uses to implement and represent the mechanisms that enforce the CIR and PIR rates.

The 7210 SAS E X uses a rate step value based on the configured rate to define the granularity for both the CIR and PIR rates (Please see the [Table 11](#) Supported Hardware Rates and CIR/PIR values for details). The adaptation rule controls the method the system uses to choose the rate step based on the administrative rates defined by the **rate** command.

Table 11: Supported Hardware Rates and CIR/PIR Values

Hardware Rate Steps	Rate Range (kbps)
12.359619	0 to 3151
24.719238	3152 to 6303
49.438477	6304 to 12606
98.876953	12607 to 25213
197.753906	25214 to 50427
395.507812	50428 to 100854
791.015625	100855 to 201708
791.015625	100855 to 201708
1582.031250	201709 to 403417
3164.062500	403418 to 806835

Table 11: Supported Hardware Rates and CIR/PIR Values (Continued)

Hardware Rate Steps	Rate Range (kbps)
6328.125000	806836 to 1613671
12656.25	1613672 to 3227343
25312.50	3227344 to 6454687
50625.0	6454688 to 12909375
101250.0	12909376 to 25818750

To illustrate how the adaptation rule constraints **minimum**, **maximum** and **closest** are evaluated in determining the operational CIR or PIR for the 7210 SAS, assume there is a queue where the administrative CIR and PIR values are 90Kbps and 150 Kbps, respectively.

If the adaptation rule is **minimum**, the operational CIR and PIR values will be 99 Kbps and 161 Kbps (The hardware rate step is 12.359619) respectively as it is the native hardware rate greater than or equal to the administrative CIR and PIR values.

If the adaptation rule is **maximum**, the operational CIR and PIR values will be 87 Kbps and 149 Kbps.

If the adaptation rule is **closest**, the operational CIR and PIR values will be 87 Kbps and 149 Kbps, respectively, as those are the closest matches for the administrative values that are even multiples of the 12.359619 Kbps rate step.

Queue Priority and Weight

The priority for the queue can be specified by using the cir-level parameter. The system maps the cir-level to a pir-level and it is not user configurable. Cir-level parameter defines the scheduling priority of the queue in the CIR loop and the PIR loop (the system assigns the priority for the queue based on its cir-level).

Cir-level value of 8, represents the highest priority. Additionally the scheduler always provides the configured bandwidth (CIR = PIR) for the queues assigned this value (if bandwidth is available), irrespective of the whether the CIR of other queues are met or not. In other words, CIR rate of level-8 queues in the system are satisfied first before satisfying the CIR rate of queues at other levels. PIR configured for queues at this level are ignored by the system.

User can specify the weight for the queue. The weight parameter is used to determine the proportion of the available bandwidth that is allocated to the queues vying for bandwidth at the same priority.

Committed Burst Size and the Maximum Burst Size (MBS)

The committed burst size and maximum burst size (CBS and MBS) parameters specify the amount of buffers reserved for a queue and upto how much of buffers a queue can contend for in the shared buffer space respectively. Once the reserved buffers for a given queue have been used, the queue contends with other queues for shared buffer resources up to the maximum burst size.

The CBS and MBS for the queues are configurable entities for the access and network ports and access uplink ports. The CBS and MBS value for the queues is set to appropriate default values which takes care of specific FC needs in terms of maintaining the differential treatment.

Service Ingress QoS Policies

Service ingress QoS policies define ingress service forwarding class and map flows to those. When a service ingress QoS policy is created, it always has two defined that cannot be deleted: one for the traffic and one for multipoint traffic. These exist within the definition of the policy. The only get instantiated in hardware when the policy is applied to a SAP. In the case where the service does not have multipoint traffic, the multipoint will not be instantiated.

In the simplest service ingress QoS policy, all traffic is treated as a single flow and mapped to a single, and all flooded traffic is treated with a single multipoint. The required elements to define a service ingress QoS policy are:

- A unique service ingress QoS policy ID.
- A QoS policy scope of template or exclusive.
- At least one default unicast forwarding class. The parameters that can be configured for a are discussed in [Meter Parameters on page 29](#).
- At least one multipoint forwarding class meter.

Optional service ingress QoS policy elements for include:

- Additional unicast meters up to a total of 8.
- Additional multipoint meters up to 31. QoS policy match criteria to map packets to a forwarding class.
- QoS policy match criteria to map packets to a forwarding class.

Each meter can have unique meter parameters to allow individual policing of the flow mapped to the forwarding class. depicts service traffic being classified into three different forwarding classes.




Figure 2: Traffic Queuing Model for Forwarding Classes

Mapping flows to forwarding classes is controlled by comparing each packet to the match criteria in the QoS policy. The ingress packet classification to forwarding class is subject to a classification policy provisioned.

[Table 12](#) lists the classification rules that are available. Only a single classification policy can be provisioned for an entity.

The IP and MAC match criteria can be very basic or quite detailed. IP and MAC match criteria are constructed from policy entries. An entry is identified by a unique, numerical entry ID. A single entry cannot contain more than one match value for each match criteria. Each match entry has an action which specifies: the forwarding class of packets that match the entry.

The entries are evaluated in numerical order based on the entry ID from the lowest to highest ID value. The first entry that matches all match criteria has its action performed. [Table 12](#) and [Table 13](#) list the supported IP and MAC match criteria.

Table 12: Service Ingress QoS Policy IP Match Criteria

IP Criteria	
DSCP value (available for SAPs in VPLS, VLL and VPRN services)	IP source and mask, IP destination and mask, IP protocol, TCP/UDP source port, TCP/UDP destination port, (available only for SAPs in VPRN services)

Table 13: Service Ingress QoS Policy MAC Match Criteria

MAC Criteria	
- IEEE 802.1p value/mask (available for SAPs in VPLS, VLL and VPRN services) - Source MAC address/mask (available for SAPs in VPLS, VLL and VPRN services) - Destination MAC address/mask (available for SAPs in VPLS, VLL and VPRN services) - EtherType value (available for SAPs in VPLS, VLL and VPRN services)	

The MAC match criteria that can be used for an Ethernet frame depends on the frame's format. See [Table 14](#).

Table 14: MAC Match Ethernet Frame Types

Frame Format	Description
802dot3	IEEE 802.3 Ethernet frame. Only the source MAC, destination MAC and IEEE 802.1p value are compared for match criteria.
Ethernet-II	Ethernet type II frame where the 802.3 length field is used as an Ethernet type (Etype) value. Etype values are two byte values greater than 0x5FF (1535 decimal).

[Table 15](#) lists the criteria that can be matched for the various MAC frame types.

Table 15: MAC Match Criteria Frame Type Dependencies

Frame Format	Source MAC	Dest MAC	IEEE 802.1p Value	Etype Value
802dot3	Yes	Yes	Yes	No
ethernet-II	Yes	Yes	Yes	Yes

Service ingress QoS policy ID 1 is reserved for the default service ingress policy. The default policy cannot be deleted or changed.

The default service ingress policy is implicitly applied to all SAPs which do not explicitly have another service ingress policy assigned. The characteristics of the default policy are listed in [Table 16](#).

Table 16: Default Service Ingress Policy ID 1 Definition

Characteristic	Item	Definition
	Meter 1	1 (one) all unicast traffic: • Forward Class: best-effort (be) • CIR = 0 • PIR = max (4000000 kbps in case of a LAG with four member ports) • MBS, CBS = default (values derived from applicable policy)
	Meter 11	1 (one) for all multipoint traffic: • CIR = 0 • PIR = max (4000000 kbps in case of a LAG with four member ports) • MBS, CBS = default (values derived from applicable policy)
Flows	Default Forwarding Class	1 (one) flow defined for all traffic: • All traffic mapped to best-effort (be)

Hierarchical Ingress Policing

Hierarchical ingress policing allows the users to specify the amount of traffic admitted into the system per SAP. It also allows the user to share the available bandwidth per SAP among the different FCs of the SAP. For example, user can allow the packets classified as Internet data to use the entire SAP bandwidth when other forwarding classes do not have traffic.

It provides an option to configure SAP aggregate policer per SAP on SAP ingress. The user should configure the PIR rate of the aggregate policer. The user can optionally configure the burst size of the aggregate policer.

The aggregate policer monitors the traffic on different FCs and determines if the packet has to be forwarded to an identified profile or dropped. The final disposition of the packet is based on the operating rate of the following:

- Per FC policer
- Per SAP aggregate policer

For more information on the final disposition of the packet, refer to the command description of "aggregate-meter-rate" command in the 7210 SAS X Services Guide.

A new meter mode "trtcm2" (RFC 4115) is introduced for use only on SAP ingress. When the SAP aggregate policer is configured, the per FC policer can be only configured in "trtcm2" mode. The existing meter mode "trtcm" is re-named as "trtcm1" (RFC 2698). The meter modes "srtCM" and "trtcm1" are used in the absence of aggregate meter.

Service Egress QoS Policies

Service egress queues are implemented at the transition from the service core network to the service access network. The advantages of per-service queuing before transmission into the access network are:

- Per-service egress subrate capabilities especially for multipoint services.
- More granular, fairer scheduling per-service into the access network.
- Per-service statistics for forwarded and discarded service packets.

The subrate capabilities and per-service scheduling control are required to make multiple services per physical port possible. Without egress shaping, it is impossible to support more than one service per port. There is no way to prevent service traffic from bursting to the available port bandwidth and starving other services.

For accounting purposes, per-service statistics can be logged. When statistics from service ingress queues are compared with service egress queues, the ability to conform to per-service QoS requirements within the service core can be measured. The service core statistics are a major asset to core provisioning tools.

Service egress QoS policies define egress queues and map forwarding class flows to queues. The system allocates 8 queues to service egress by default. To define a basic egress QoS policy, the following are required:

- A unique service egress QoS policy ID.
- A QoS policy scope of template or exclusive.
- The parameters that can be configured for a queue are discussed in [Queue Parameters on page 35](#).

Optional service egress QoS policy elements include:

- Specify remark policy that defines IEEE 802.1p priority value remarking based on forwarding class.

In 7210 SAS-X, the '**sap-qos-marking**' command is provided that allows option for the user to configure SAP-based marking and port-based marking. In SAP-based marking, the remark policy defined in the SAP egress policy associated with each SAP is used to mark the packets egressing out of SAP if marking is enabled. In port-based marking, the remark policy defined in the access-egress policy associated with the access port determines the marking values to use for all the SAPs defined on that port. More information is available in the section on Access Egress policies in this guide.

Each queue in a policy is associated with one of the forwarding classes. Each queue can have its individual queue parameters allowing individual rate shaping of the forwarding class(es) mapped to the queue.

The forwarding class determination per service egress packet is determined at ingress. If the packet ingressed the service on the same 7210 SAS X router, the service ingress classification rules determine the forwarding class of the packet. If the packet is received, the forwarding class is marked in the tunnel transport encapsulation.

Service egress QoS policy ID 1 is reserved as the default service egress policy. The default policy cannot be deleted or changed. The default egress policy is applied to all SAPs which do not have another service egress policy explicitly assigned. The characteristics of the default policy are listed in the following table.

Table 17: Default SAP Egress Policy ID 1 Definition

Characteristic	Item	Definition
Queues	Queue 1-8	1 (one) queue defined for each traffic class
Network-Control (nc)	Queue 8	• CIR=0 • PIR=max (line rate) • Cir-Level 1 • Pir-Weight 1 • Queue-Management Policy : default
High-1 (h1)	Queue7	• CIR=0 • PIR=max (line rate) • Cir-Level 1 • Pir-Weight 1 • Queue-Management Policy : default
Expedited (ef)	Queue 6	• CIR = 0 • PIR = max (line rate) • Cir-Level 1 • Pir-Weight 1 • Queue-Management Policy : default
High-2 (h2)	Queue 5	• CIR = 0 • PIR = max (line rate) • Cir-Level 1 • Pir-Weight 1 • Queue-Management Policy : default
Low-1 (l1)	Queue 4	• CIR = 0 • PIR = max (line rate) • Cir-Level 1 • Pir-Weight 1 • Queue-Management Policy : default
Assured (af)	Queue 3	• CIR = 0 • PIR = max (line rate)

Table 17: Default SAP Egress Policy ID 1 Definition (Continued)

Characteristic	Item	Definition
Low-2 (l2)	Queue 2	• Cir-Level 1
		• Pir-Weight 1
		• Queue-Management Policy : default
		• CIR = 0
		• PIR = max (line rate)
		• Cir-Level 1
Best-Effort (be)	Queue 1	• Pir-Weight 1
		• Queue-Management Policy : default
		• CIR = 0
		• PIR = max (line rate)
		• Cir-Level 1
		• Pir-Weight 1
Flows	Default Action	• Queue-Management Policy : default
		All FCs are mapped to corresponding Queues and Dot1p values are marked as follows:

Table 17: Default SAP Egress Policy ID 1 Definition (Continued)

Characteristic	Item	Definition	
		In-Profile	Out-Profile
Network-Control (nc)		7	7
High-1(h1)		6	6
Expedited (ef)		5	5
High-2 (h2)		4	4
Low-1 (l1)		3	3
Assured (af)		2	2
Low-2 (l2)		1	1
Best-Effort (be)		0	0

Access Egress QoS Policies

An access egress policy defines marking values for the traffic egressing towards the customer on the access ports. Access egress QoS policies map forwarding class flows to marking values to use.

7210 SAS-X supports SAP-based egress marking and port-based egress marking on access ports. Users have an option to turn on either sap-based marking or port-based marking using the command 'sap-qos-marking' under the configure>port>ethernet>access>egress context. In SAP-based marking the remark policy defined in the SAP egress policy associated with each SAP is used to mark the packets egressing out of SAP if marking is enabled. In port-based marking, the remark policy defined in the access-egress policy associated with the access port determines the marking values to use for all the SAPs defined on that port. SAP-based marking is only supported for L2 SAPs, i.e. SAPs configured in Epipe, VPLS and PBB (I-SAPs only) service. Port-based marking is supported for L3 SAPs (i.e. SAPs configured in VPRN services), PBB B-SAPs and other L2 SAPs configured on the port. The access egress policy is used only when port-based marking has been enabled (i.e. sap-qos-marking is set to disable). More information on the CLI command 'sap-qos-marking' is available in the 7210 SAS Interfaces guide.

A remarking policy can be defined for each access egress policy and remarking is disabled by default. Only remarking policy of type dot1p, dot1p-lsp-exp-shared, dscp or dot1p-dscp can be used with access-egress policy.

To define a basic access egress QoS policy, the following are required:

- A unique service access QoS policy ID.
- A QoS policy scope of template or exclusive.
- A remark policy of appropriate type for remarking based on forwarding class.

Remarking by default is disabled. It can be enabled by the remarking command present under access egress context.

Access egress QoS policy ID 1 is reserved as the default access egress policy. The default policy cannot be deleted or changed. The default access egress policy is applied to all access ports which do not have another access egress policy explicitly assigned. By default sap-qos-marking is enabled. The default access-egress policy is as shown below

```
-----
*A:Dut-A>config>qos>access-egress# info detail
-----
      description "Default Access egress QoS policy."
      no remarking
      remark 2
-----
```

Buffer Pools

The 7210 SAS X has a single buffer pool per node, the system pool. All the queues created by the system are allocated buffers from this system pool. Queues come up with default buffers, and the buffers change accordingly when they are associated with a network port or SAP. Queue management policies allow the user to specify the parameters that determine buffer allocation to the queues.

Queue Management Policies

,Queue management policies allows the user to define the queue buffer and WRED slope parameters.

The 7210 SAS supports a single buffer pool per node. All the queues created in the system are allocated buffers from this system pool. The default buffers are allocated to the queues accordingly when they are associated with a SAP or a network port.

Queue management policies allow the user to define the CBS, MBS and WRED parameters for use by the queue. The CBS and MBS parameters are used to allocate the appropriate amount of buffers from the system pool to the queues. The WRED parameters allow the user to define the WRED slope characteristics. User can define a high-slope and a low-slope for each of the queues. High-slope is used for in-profile packets being enqueued into the queues and low-priority slope is used for out-of-profile packets being enqueued into the queues.

By default each queue is associated with a default queue-management policy. The default policy allocates the appropriate amount of CBS and MBS buffers based on whether the queue is associated with a SAP or network port.

WRED Slopes

Operation and Configuration

The 7210 SAS provides a single system buffer pool for use by all the queues created in the system. Each queue supports a high-priority WRED slope, and a low-priority WRED slope. The high-priority WRED slope manages access to the shared portion of the buffer pool for high-priority or in-profile packets. The low-priority WRED slope manages access to the shared portion of the buffer pool for low-priority or out-of-profile packets.

By default, the high-priority and low-priority slopes are disabled.

In the 7210 SAS X, WRED is supported. WRED uses average queue lengths, queue thresholds provisioned, and drop probability to calculate the packet's eligibility to be enqueued. The committed portion of the buffer pool is exclusively used by a queue to enqueue traffic within committed rate.

In the 7210 SAS X, WRED is supported. WRED uses average queue lengths, queue thresholds provisioned, and drop probability to calculate the packet's eligibility to be enqueued. The committed portion of the buffer pool is exclusively used by a queue to enqueue traffic within committed rate.

For the queues within a buffer pool, packets are either queued using committed burst size (CBS) buffers or shared buffers. The CBS buffers are simply buffer memory that has been allocated to the queue while the queue depth is at or below its CBS threshold. The amount of CBS assigned to all queues is dependent upon the number of queues created, the setting of the default CBS as defined in the policy, and any CBS values set per queue within a QoS policy. However, from a functional perspective, the buffer pool does not keep track of the total of the CBS assigned to queues serviced by the pool. CBS subscription on the pool is an administrative function that must be monitored by the queue provisioner.

For each queue, the amount of access and network buffer pools, the percentage of the buffers that are to be reserved for CBS buffers is configured by the usersoftware (cannot be changed by user). This setting indirectly assigns the amount of shared buffers on the pool. This is an important function that controls the ultimate average and total shared buffer utilization value calculation used for WRED slope operation. The CBS setting can be used to dynamically maintain the buffer space on which the WRED slopes operate.

When a queue depth exceeds the queue's CBS, packets received on that queue must contend with other queues exceeding their CBS for shared buffers. To resolve this contention, the buffer pool uses two WRED slopes to determine buffer availability on a packet by packet basis. A packet that was either classified as high priority or considered in-profile is handled by the high-priority WRED slope. This slope should be configured with WRED parameters that prioritize buffer

availability over packets associated with the low-priority WRED slope. Packets that had been classified as low priority or out-of-profile are handled by this low-priority WRED slope.

The following is a simplified overview of how a WRED slope determines shared buffer availability on a packet basis:

1. The WRED function keeps track of shared buffer utilization and shared buffer average utilization.
2. At initialization, the utilization is 0 (zero) and the average utilization is 0 (zero).
3. When each packet is received, the current average utilization is plotted on the slope to determine the packet's discard probability.
4. A random number is generated associated with the packet and is compared to the discard probability.
5. The lower the discard probability, the lower the chances are that the random number is within the discard range.
6. If the random number is within the range, the packet is discarded which results in no change to the utilization or average utilization of the shared buffers.
7. A packet is discarded if the utilization variable is equal to the shared buffer size or if the utilized CBS (actually in use by queues, not just defined by the CBS) is oversubscribed and has stolen buffers from the shared size, lowering the effective shared buffer size equal to the shared buffer utilization size.
8. If the packet is queued, a new shared buffer average utilization is calculated using the time-average-factor (TAF) for the buffer pool. The TAF describes the weighting between the previous shared buffer average utilization result and the new shared buffer utilization in determining the new shared buffer average utilization. (See [Tuning the Shared Buffer Utilization Calculation on page 56.](#))
9. The new shared buffer average utilization is used as the shared buffer average utilization next time a packet's probability is plotted on the WRED slope.
10. When a packet is removed from a queue (if the buffers returned to the buffer pool are from the shared buffers), the shared buffer utilization is reduced by the amount of buffers returned. If the buffers are from the CBS portion of the queue, the returned buffers do not result in a change in the shared buffer utilization.




Figure 3: RED Slope Characteristics

A RED slope itself is a graph with an X (horizontal) and Y (vertical) axis. The X-axis plots the percentage of shared buffer average utilization, going from 0 to 100 percent. The Y-axis plots the probability of packet discard marked as 0 to 1. The actual slope can be defined as four sections in (X, Y) points (Figure 3):

1. Section A is (0, 0) to (start-avg, 0). This is the part of the slope that the packet discard value is always zero, preventing the RED function from discarding packets when the shared buffer average utilization falls between 0 and start-avg.
2. Section B is (start-avg, 0) to (max-avg, max-prob). This part of the slope describes a linear slope where packet discard probability increases from zero to max-prob.
3. Section C is (max-avg, max-prob) to (max-avg, 1). This part of the slope describes the instantaneous increase of packet discard probability from max-prob to one. A packet discard probability of 1 results in an automatic discard of the packet.
4. Section D is (max-avg, 1) to (100%, 1). On this part of the slope, the shared buffer average utilization value of max-avg to 100% results in a packet discard probability of 1.

Plotting any value of shared buffer average utilization will result in a value for packet discard probability from 0 to 1. Changing the values for start-avg, max-avg and max-prob allows the adaptation of the RED slope to the needs of the access or network queues using the shared portion of the buffer pool, including disabling the RED slope.

Tuning the Shared Buffer Utilization Calculation

The 7210 SAS X allows tuning the calculation of the Shared Buffer Average Utilization (SBAU) after assigning buffers for a packet entering a queue as used by the RED slopes to calculate a packet's drop probability. The 7210 SAS X implements a time average factor (TAF) parameter in the buffer policy which determines the contribution of the historical shared buffer utilization and the instantaneous Shared Buffer Utilization (SBU) in calculating the SBAU. The TAF defines a

weighting exponent used to determine the portion of the shared buffer instantaneous utilization and the previous shared buffer average utilization used to calculate the new shared buffer average utilization. To derive the new shared buffer average utilization, the buffer pool takes a portion of the previous shared buffer average and adds it to the inverse portion of the instantaneous shared buffer utilization (SBU). The formula used to calculate the average shared buffer utilization is:

$$SBAU_n = \left(SBU \times \frac{1}{2^{TAF}} \right) + \left(SBAU_{n-1} \times \frac{2^{TAF} - 1}{2^{TAF}} \right)$$

where:

$SBAU_n$ = Shared buffer average utilization for event n

$SBAU_{n-1}$ = Shared buffer average utilization for event (n-1)

SBU = The instantaneous shared buffer utilization

TAF = The time average factor

Table 18 shows the effect the allowed values of TAF have on the relative weighting of the instantaneous SBU and the previous SBAU ($SBAU_{n-1}$) has on the calculating the current SBAU ($SBAU_n$).

Table 18: TAF Impact on Shared Buffer Average Utilization Calculation

TAF	2^{TAF}	Equates To	Shared Buffer Instantaneous Utilization Portion	Shared Buffer Average Utilization Portion
0	2^0	1	1/1 (1)	0 (0)
1	2^1	2	1/2 (0.5)	1/2 (0.5)
2	2^2	4	1/4 (0.25)	3/4 (0.75)
3	2^3	8	1/8 (0.125)	7/8 (0.875)
4	2^4	16	1/16 (0.0625)	15/16 (0.9375)
5	2^5	32	1/32 (0.03125)	31/32 (0.96875)
6	2^6	64	1/64 (0.015625)	63/64 (0.984375)
7	2^7	128	1/128 (0.0078125)	127/128 (0.9921875)
8	2^8	256	1/256 (0.00390625)	255/256 (0.99609375)

Table 18: TAF Impact on Shared Buffer Average Utilization Calculation (Continued)

TAF	2^{TAF}	Equates To	Shared Buffer Instantaneous Utilization Portion	Shared Buffer Average Utilization Portion
9	2^9	512	1/512 (0.001953125)	511/512 (0.998046875)
10	2^{10}	1024	1/1024 (0.0009765625)	1023/2024 (0.9990234375)
11	2^{11}	2048	1/2048 (0.00048828125)	2047/2048 (0.99951171875)
12	2^{12}	4096	1/4096 (0.000244140625)	4095/4096 (0.999755859375)
13	2^{13}	8192	1/8192 (0.0001220703125)	8191/8192 (0.9998779296875)
14	2^{14}	16384	1/16384 (0.00006103515625)	16383/16384 (0.99993896484375)
15	2^{15}	32768	1/32768 (0.000030517578125)	32767/32768 (0.999969482421875)

The value specified for the TAF affects the speed at which the shared buffer average utilization tracks the instantaneous shared buffer utilization. A low value weights the new shared buffer average utilization calculation more to the shared buffer instantaneous utilization. When TAF is zero, the shared buffer average utilization is equal to the instantaneous shared buffer utilization. A high value weights the new shared buffer average utilization calculation more to the previous shared buffer average utilization value. The TAF value applies to all high and low priority RED slopes for ingress and egress buffer pools controlled by the buffer policy.

Queue Management Policy Parameters

The elements required to define a queue management policy are:

- A unique policy ID
- The high and low RED slope shapes for the queue: the start-avg, max-avg and maxprob settings for the high-priority and low-priority RED slopes.
- The TAF weighting factor to use for the SBAU calculation for determining RED slope drop probability.

Queue management policy ID **default** is reserved for the default queue management policy. The default policy cannot be deleted or changed. The default policy is implicitly applied to all queues which do not have another queue management policy explicitly assigned.

[Table 19](#) lists the default values for the default slope policy.

Table 19: Default Slope Policy Definition

Parameter	Description	Setting
Policy ID	Queue management policy ID	default (for default queue management policy)
CBS	Committed Burst size	Default (in kilobytes)
MBS	Maximum Burst size	Default (in kilobytes)
High (RED) slope	Administrative state	Shutdown
	start-avg	70% utilization
	max-avg	90% utilization
	max-prob	75%
Low (RED) slope	Administrative state	Shutdown
	start-avg	50% utilization
	max-avg	75% utilization
	max-prob	75%

Remark Policy

This policy allows the user to define the forwarding class to egress marking values. Based on the packet encapsulation used to send out the service packets, the remark policy allows the user to define and associate appropriate policies to service egress and network egress QoS policies. 7210 supports the following types of remark policies:

- dot1p (for use in service egress and network qos [port type] policies)
- dscp (for use in network qos [port type] policies)
- lsp-exp (for use in network qos [ip-interface type] policies)
- dot1p-dscp (for use in network qos [port type] policies)
- dot1p-lsp-exp-shared (for use in service egress and network qos [ip-interface type] policies)

Each of these remark policy type can be associated with only appropriate Qos policies as listed above.

The required elements to define a remark QoS policy are:

- A unique remark QoS policy ID.
- Forwarding class to appropriate marking values

Table 20: Slope Policy Defaults

Field	Default
description	Default slope policy
high (RED) slope	
Administrative state	shutdown
start-threshold	75% utilization
queue 1 — 8 drop-rate	1 (6.25% drop rate)
low (RED) slope	
Administrative state	shutdown
start-threshold	50% utilization
queue 1 — 8 drop-rate	0 (100% drop rate)

Egress Port Rate Limiting

The 7210 SAS supports port egress rate limiting. This feature allows the user to limit the bandwidth available on the egress of the port to a value less than the maximum possible link bandwidth. It also allows the user to control the amount of burst sent out.

Forwarding Classes

7210 SAS support multiple forwarding classes and class-based queuing, so the concept of forwarding classes is common to all of the QoS policies.

Each forwarding class (also called Class of Service (CoS)) is important only in relation to the other forwarding classes. A forwarding class provides network elements a method to weigh the relative importance of one packet over another in a different forwarding class.

Queues are created for a specific forwarding class to determine the manner in which the queue output is scheduled. The forwarding class of the packet, along with the in-profile or out-of-profile state, determines how the packet is queued and handled (the per hop behavior (PHB)) at each hop along its path to a destination egress point. 7210 SAS support eight (8) forwarding classes ([Table 21](#)).

Table 21: Forwarding Classes

FC-ID	FC Name	FC Designation	DiffServ Name	Notes
7	Network Control	NC	NC2	Intended for network control traffic.
6	High-1	H1	NC1	Intended for a second network control class or delay/jitter sensitive traffic.
5	Expedited	EF	EF	Intended for delay/jitter sensitive traffic.
4	High-2	H2	AF4	Intended for delay/jitter sensitive traffic.
3	Low-1	L1	AF2	Intended for assured traffic. Also is the default priority for network management traffic.
2	Assured	AF	AF1	Intended for assured traffic.
1	Low-2	L2	CS1	Intended for BE traffic.
0	Best Effort	BE	BE	

Note that [Table 21](#) presents the default definitions for the forwarding classes. The forwarding class behavior, in terms of ingress marking interpretation and egress marking, can be changed by a [Network QoS Policies on page 23](#). All forwarding class queues support the concept of in-profile and out-of-profile.

Forwarding-Class To Queue-ID Map

There are 8 forwarding classes supported on 7210 SAS X. Each of these FC is mapped to a specific queue while traffic is flowing on the egress port. By mapping FC to different queues the differential treatment is imparted to various classes of traffic.

The 7210 SAS allocates 8 queues to SAP and network by default. The queues cannot be created or deleted by the user. CLI commands are available for the user to configure the queue parameters.

The queue ID 1 to 8 are assigned to each of the 8 queues. Queue-ID 8 is the highest priority and queue-id 1 is the lowest priority. FCs are correspondingly mapped to these queue IDs according to their priority. The system defined map is as shown in [Table 22](#). This mapping is not user configurable.

Table 22: Forwarding Class to Queue-ID Map

FC-ID	FC Name	FC Designation	Queue-ID
7	Network control	NC	8
6	High-1	H1	7
5	Expedited	EF	6
4	High-2	H2	5
3	Low-1	L1	4
2	Assured	AF	3
1	Low-2	L2	2
0	Best-Effort	BE	1

QoS Policy Entities

Services are configured with default QoS policies. Additional policies must be explicitly created and associated. There is one default service ingress QoS policy, one default service egress QoS policy, one default network QoS policy and default network queue policy. Only one ingress QoS policy and one egress QoS policy can be applied to a SAP or IP interface, or network port.

When you create a new QoS policy, default values are provided for most parameters with the exception of the policy ID, descriptions. Each policy has a scope, default action, a description, and meters for ingress policies and queues for egress policies. The queue is associated with a forwarding class.

QoS policies can be applied to the following service types:

- Epipe — SAP ingress and egress policies are supported on an Epipe service access point (SAP).
- VPLS — SAP ingress and egress policies are supported on a VPLS SAP.
- VPRN — Only ingress policies are supported on a VPRN SAP.

QoS policies can be applied to the following entities:

- SAP ingress and egress policies on access SAPs.
- Network QoS Policy (ingress) on Network Port and/or Network IP Interface.
- Network Queue Policy on Network port.
- Multipoint bandwidth management policies to manage multipoint traffic (per system).

Default Access QoS policies maps all traffic with equal priority and allow an equal chance of transmission (Best Effort (be) forwarding class) and an equal chance of being dropped during periods of congestion. QoS prioritizes traffic according to the forwarding class and uses congestion management to control access ingress, access egress, and network traffic with queuing according to priority.

Configuration Notes

The following information describes QoS implementation caveats:

- Creating additional QoS policies is optional.
- Default policies are created for service ingress, service egress, network-ingress, network-egress, queue-management and multipoint bandwidth management.
- Associating a service or ports with a QoS policy other than the default policy is optional.

Port Level Egress Rate-Limiting

In This Section

This section provides information to configure port level egress-rate using the command line interface.

Topics in this section include:

- [Overview on page 68](#)
- [Basic Configurations on page 70](#)
- [Configuration Commands on page 74](#)

Overview

Egress port rate limiting allows the device to limit the traffic that egresses through a port to a value less than the available link bandwidth. This feature is supported on the 7210 SAS-Series platforms.

Applications

This feature is useful when connecting the 7210 SAS to an Ethernet-over-SDH (EoSDH) (or microwave) network, where the network allocates predetermined bandwidth to the nodes connecting into it, based on the transport bandwidth requirement. When connecting to such a network it is important that the traffic sent into the SDH node does not exceed the configured values, since the SDH network does not have QoS capabilities and buffers required to prioritize the ingress traffic.

Egress rate attributes include:

- Allows for per port configuration of the maximum egress port rate, using the egress-rate CLI command.
- Ethernet ports support this feature.
- The scheduler distributes the available maximum egress bandwidth based on the CIR/PIR configuration parameters provisioned for the queues.
- Provides support for a burst parameter to control the amount of burst the egress port can generate.
- When ports are members of a LAG, all the ports use the same value for the egress-rate and the max-burst parameters.
- If frame overhead accounting is enabled, then queue scheduler accounts for the Ethernet frame overhead.

Affect of Port Level Rate-Limiting on Queue Functionality

- When an egress-rate sub-rate value is given, the queue rates that are specified using percentages will use the egress-rate value instead of the port bandwidth to configure the appropriate queue rates. Configuration of egress port rate to different values will result in a corresponding dynamic adjustment of rates for the queues configured on ports.
- When the egress-rate sub-rate value is set, CBS/MBS of the associated network queues will not change.

Basic Configurations

To apply port level rate-limiting, perform the following:

- The **egress-rate** command is present in the ***A:Dut-1>config>port>ethernet** context.
- The **egress-rate** configures the maximum rate (in kbps) for the port. The value should be between 1 and 1000000 kbps and between 1 and 10000000 kbps for 10G port.
- The **max-burst** command configures a maximum-burst (in kilo-bits) associated with the egress-rate. This is optional parameter and if not defined then, by default, it is set to 32kb for a 1G port and 66kb for a 10G port. User cannot configure max-burst without configuring egress-rate. The value should be between 32 and 16384 or default.
- By default there is no egress-rate command set on port. By default egress-rate for a port is maximum (equal to line-rate).

The following displays the egress-rate configuration for a port:

```
*A:Dut-1>config>port# info
-----
    ethernet
        egress-rate 120000 max-burst 234
    exit
    no shutdown
-----
*A:Dut-1>config>port#
```

Modifying Port Level Egress-Rate Command

To modify egress-rate parameters you can simply apply a egress-rate command with new egress-rate and max-burst value.

The following displays the egress-rate configuration for a port:

```
*A:Dut-1>config>port# ethernet egress-rate 10000 max-burst default
*A:Dut-1>config>port# info
-----
    ethernet
        egress-rate 10000
    exit
    no shutdown
-----
*A:Dut-1>config>port#
```

Removing Port Level Egress-Rate Command

To remove egress-rate command from a port, use the **no** option with egress-rate command. The rate for egress-rate option and the max-burst should not be used in this case.

CLI Syntax: `config>port>ethernet# no egress-rate`

The following displays the removal of egress-rate configuration from a port:

```
*A:Dut-1>config>port# no ethernet egress-rate
*A:Dut-1>config>port# info
-----
        ethernet
        exit
        no shutdown
-----
*A:Dut-1>config>port#
```

Default Egress-Rate Values

By default no egress-rate is configured for a port.

Port Level Egress-Rate Command Reference

Command Hierarchies

Configuration Commands

```
config
  — port
    — ethernet
      — egress-rate <sub-rate> [max-burst <size-in-kbits>]
      — no egress-rate
```

Show Commands

```
show
  — port [port-id]
```

Configuration Commands

egress-rate

Syntax	egress-rate <sub-rate> [max-burst <size-in-kbits>] no egress-rate
Context	config>port>ethernet
Description	This command configures maximum rate and corresponding burst value for a port. The egress-rate is configured as kbps while max-burst is configured as kilo-bits while max-burst should be between 32 and 16384 or default. The no form of the command removes egress-rate from the port.
Default	No egress-rate and max-burst is configured for the port.
Parameters	<i>sub-rate</i> — Specifies an integer value between 1 and 1000000 kbps and between 1 and 10000000 kbps for 10G port. max-burst — Specifies an integer value, in kilo-bits, between 32 and 16384 or default.

Show Commands

port

Syntax	port [<i>port-id</i>]
Context	show
Description	This command displays Egress-Rate and Max-Burst value set for port along with other details of the port.
Parameters	<i>port-id</i> — Displays information about the specific port ID.

Sample Output

```
*A:Dut-1>config>port>ethernet# show port 1/1/23
=====
Ethernet Interface
=====
Description      : 10/100/Gig Ethernet SFP
Interface        : 1/1/23                      Oper Speed      : 100 mbps
Link-level       : Ethernet                    Config Speed    : 1 Gbps
Admin State      : up                          Oper Duplex     : full
Oper State       : up                          Config Duplex   : full
Physical Link    : Yes                         MTU             : 9212
IfIndex          : 36405248                     Hold time up    : 0 seconds
Last State Change : 03/12/2001 03:31:09          Hold time down  : 0 seconds
Last Cleared Time : N/A

Configured Mode   : network                      Encap Type      : null
Dot1Q Ethertype  : 0x8100                       QinQ Ethertype  : 0x8100
Net. Egr. Queue Pol: default                     Access Egr. Qos *: n/a
Egr. Sched. Pol  : default                       Network Qos Pol : 1
Auto-negotiate   : true                          MDI/MDX        : MDX
Accounting Policy : None                         Collect-stats   : Disabled
Egress Rate      : 100000                         Max Burst      : 8000

Down-when-looped : Disabled                       Keep-alive      : 10
Loop Detected    : False                          Retry           : 120

Configured Address : 00:f7:d6:5e:98:18
Hardware Address   : 00:f7:d6:5e:98:18
Cfg Alarm          :
Alarm Status       :

Transceiver Data

Transceiver Type   : SFP
Model Number       : 3HE00062AAAA01 ALA IPUIAEHDAA6
TX Laser Wavelength: 0 nm                         Diag Capable    : no
Connector Code     : Unknown                       Vendor OUI      : 00:90:65
```

Port Level Egress-Rate Command Reference

```

Manufacture date   : 2008/09/11           Media           : Ethernet
Serial Number      : PEB2WGH
Part Number        : FCMJ-8521-3-A5
Optical Compliance : GIGE-T
Link Length support: 100m for copper
  
```

Traffic Statistics

```

=====
                                     Input           Output
-----
Octets                               15028477        3236
Packets                             16729           19
Errors                               0              0
=====
  
```

* indicates that the corresponding row element may have been truncated.

Port Statistics

```

=====
                                     Input           Output
-----
Unicast Packets                      11611          17
Multicast Packets                     359            0
Broadcast Packets                     4759            2
Discards                             0              0
Unknown Proto Discards                0
=====
  
```

Ethernet-like Medium Statistics

```

=====
Alignment Errors : 0  Sngl Collisions : 0
FCS Errors       : 0  Mult Collisions : 0
SQE Test Errors  : 0  Late Collisions : 0
CSE              : 0  Excess Collisns : 0
Too long Frames  : 0  Int MAC Tx Errs  : 0
Symbol Errors    : 0  Int MAC Rx Errs  : 0
=====
  
```

*A:MTU-T2>config>port>ethernet#

Frame Based Accounting

In This Section

This section provides information to configure frame-based accounting using the command line interface.

Topics in this section include:

- [Overview on page 78](#)
- [Basic Configurations on page 79](#)
- [Configuration Commands on page 82](#)

Overview

This feature when enabled let QoS policies to accounts for the Ethernet frame overhead (for example, it accounts for the IFG (inter-frame gap) and the preamble). Typically, the IFG and preamble constitutes about $12 + 8 = 20$ bytes. The overhead for Ethernet ports uses this value.

A configurable CLI command enables accounting of the frame overhead at ingress. This is a system wide parameter and affects the behavior of the ingress meter. By default frame-based accounting is disabled on ingress. Frame overhead is always accounted for at the egress (queue rates and egress rate) and user has no option of disabling it.

Affects of Enabling Ingress Frame Based Accounting on Ingress Meter Functionality

To enable system-wide consistency in configuring QoS queue and meter rate parameters, the meters used on the system ingress might need to account for Ethernet frame overhead. ingress and service ingress meters account for Ethernet frame overhead. A configurable CLI command can enable or disable the frame overhead accounting. This is a system-wide parameter affecting the behavior of all the meters in the system.

Affects of Egress Frame Based Accounting on Queue Functionality

Because of frame overhead accounting consideration, queue scheduler accounts for the Ethernet frame overhead. The maximum egress bandwidth accounts for the Ethernet frame overhead (it accounts for the IFG (inter-frame gap) and the preamble). Typically, the IFG and preamble constitutes about $12 + 8 = 20$ bytes. The overhead for Ethernet ports uses this value.

Accounting and Statistics

Accounting logs and statistics do not account for frame overhead.

Basic Configurations

To enable frame-based accounting, you must perform the following:

- The **frame-based-accounting** command is in the ***A:Dut-1> config>qos>frame-based-accounting** context.
- The **ingress-enable** command enables frame-based-accounting for ingress metering.

The following displays the frame-based accounting configuration:

```
*A:Dut-1>config>qos>frame-based-accounting# info detail
-----
no ingress-enable
*A:Dut-1>config>qos>frame-based-accounting#
```

Enabling and Disabling Frame-Based Accounting

To enable frame-based-accounting for ingress, you can simply use the **ingress-enable** command .
To disable frame-based-accounting for ingress, execute the **no ingress-enable** command.

CLI Syntax: `config>qos>frame-based-accounting`

The following output displays the enabling of frame-based-accounting:

```
*A:Dut-1>config>qos>frame-based-accounting# ingress-enable
*A:A:Dut-1>config>qos>frame-based-accounting# info
-----
ingress-enable
*A:A:Dut-1>config>qos>frame-based-accounting#
```

The following output displays the disabling of frame-based-accounting:

```
*A:A:Dut-1>config>qos>frame-based-accounting# no ingress-enable
*A:A:Dut-1>config>qos>frame-based-accounting#
*A:A:Dut-1>config>qos>frame-based-accounting# info detail
-----
no ingress-enable
*A:A:Dut-1>config>qos>frame-based-accounting#
```

Default Frame-Based-Accounting Values

By default frame-based-accounting is disabled for ingress. By default frame-based-accounting is enabled for egress and it cannot be disabled.

Frame Based Accounting Command Reference

Command Hierarchies

Configuration Commands

```
config
  — qos
    — frame-based-accounting
      — [no] ingress-enable
```

Show Commands

```
show
  — qos
    — network [ policy-id] [detail]
    — network-queue [network-queue-policy-name] [detail]
    — sap-egress [ policy-id] [association|detail]
    — sap-ingress [policy-id][association|match-criteria|detail]
```

Configuration Commands

ingress-enable

Syntax [no] ingress-enable

Context config>qos>frame-based-accounting

Description This command enables the frame-based-accounting for sap-ingress and network QoS.

The **no** form of the command disables frame-based-accounting for sap-ingress and network QoS.

Default disabled

Show Commands

sap-ingress

Syntax	sap-ingress [<i>policy-id</i>][association match-criteria detail]
Context	show>qos
Description	This command displays accounting status of a sap-ingress policy along with other details of the policy. When frame-based-accounting is enabled accounting is shown as frame-based otherwise packet-based.
Parameters	<i>policy-id</i> — Displays information about the specific policy ID. associations — Displays the associations of the sap-ingress policy. match-criteria — Displays the match criteria of the sap-ingress policy. detail — Displays the detailed information of the sap-ingress policy.

Sample Output

```
*A:Dut-1# show qos sap-ingress 1
=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (1)
-----
Policy-id           : 1                      Scope           : Template
Default FC         : be
Criteria-type      : None
Accounting         : frame-based
Classifiers Allowed: 16                      Meters Allowed   : 8
Classifiers Used   : 2                      Meters Used      : 2
Description        : Default SAP ingress QoS policy.
=====
*A:Dut-1#
```

network

Syntax	network [<i>policy-id</i>] [detail]
Context	show>qos
Description	This command displays the accounting status of a network qos policy along with other details of the policy. When frame-based-accounting is enabled accounting is shown as frame-based otherwise packet-based.

Show Commands

Parameters *policy-id* — Displays information about the specific policy ID.
 detail — Displays the detail policy information.

Sample Output

```
*A:Dut-1# show qos network 1
=====
QoS Network Policy
=====
-----
Network Policy (1)
-----
Policy-id      : 1                      Remark      : False
Forward Class  : be                     Profile      : Out
Attach Mode    : l2                     Config Mode  : l2+mpls
Scope          : Template               Policy Type  : port
Accounting     : frame-based
Description    : Default network-port QoS policy.
-----
Meter Mode    CIR Admin  CIR Rule  PIR Admin  PIR Rule  CBS      MBS
-----
1      TrTcm_CA  0          closest    max       closest   32       128
-----
FC                      UCastM      MCastM
-----
No FC-Map Entries Found.
-----
Port Attachments
-----
Port-id : 1/1/3
Port-id : 1/1/6
Port-id : 1/1/7
Port-id : 1/1/8
Port-id : 1/1/9
Port-id : 1/1/10
Port-id : 1/1/11
Port-id : 1/1/12
Port-id : 1/1/13
Port-id : 1/1/14
Port-id : 1/1/15
Port-id : 1/1/16
Port-id : 1/1/17
Port-id : 1/1/18
Port-id : 1/1/19
Port-id : 1/1/20
Port-id : 1/1/21
Port-id : 1/1/22
Port-id : 1/1/23
Port-id : 1/1/24
=====
*A:Dut-1#
```

sap-egress

Syntax	sap-egress [<i>policy-id</i>][association] detail
Context	show>qos
Description	This command displays accounting status of an sap-egress policy along with other details of the policy. When frame-based-accounting is enabled accounting is shown as frame-based otherwise packet-based.
Parameters	<i>policy-id</i> — Displays information about the specific policy ID. association — Displays the policy associations. detail — Displays the policy information in detail.

Sample Output

```
*A:SAS-X-C# show qos sap-egress 1
=====
QoS Sap Egress
=====
-----
Sap Egress Policy (1)
-----
Scope      : Template
Remark    : False   Remark Pol Id : 2
Accounting : frame-based
Description : Default SAP egress QoS policy.
=====
```

network-queue

Syntax	network-queue [<i>network-queue-policy-name</i>] [detail]
Context	show>qos
Description	This command displays accounting status of a network-queue policy along with other details of the policy. When frame-based-accounting is enabled accounting is shown as frame-based otherwise packet-based.
Parameters	<i>network-queue-policy-name</i> — Displays information about the specific Network queue policy. detail — Displays the detailed policy information.

Sample Output

```
*A:Dut-1# show qos network-queue default
=====
```

Show Commands

```
QoS Network Queue Policy
=====
-----
Network Queue Policy (default)
-----
Policy          : default
Accounting      : frame-based
Description     : Default network queue QoS policy.
-----
Associations
-----
Port-id : 1/1/6
Port-id : 1/1/7
Port-id : 1/1/8
Port-id : 1/1/9
Port-id : 1/1/10
Port-id : 1/1/11
Port-id : 1/1/12
Port-id : 1/1/13
Port-id : 1/1/14
Port-id : 1/1/15
Port-id : 1/1/16
Port-id : 1/1/17
Port-id : 1/1/18
Port-id : 1/1/20
Port-id : 1/1/21
Port-id : 1/1/22
Port-id : 1/1/23
Port-id : 1/1/24
=====
*A:Dut-1#
```

Network QoS Policies

In This Section

This section provides information to configure network QoS policies using the command line interface.

Topics in this section include:

- [Overview on page 88](#)
- [Basic Configurations on page 96](#)
- [Default Network Policy Values on page 99](#)
- [Service Management Tasks on page 103](#)

Overview

The network QoS policy consists of an ingress and egress component. There are two types of network QoS policies, network QoS policy of type **port** and network QoS policy of type **ip-interface**. A **port** network policy is applied to network ports, used for classification/remarking of IP traffic using DSCP or Dot1p values. Either DSCP or Dot1p can be used for ingress classification but not both. Both DSCP and Dot1p can be configured at egress for remarking. The **ip-interface** type network policy is applied to IP Interface, used for classification/remarking of MPLS traffic using EXP values. Note that the FC to Dot1p marking values configured on the port, is also used to mark the Dot1p in the VLAN tag, if any, used for MPLS traffic.

The ingress component of the policy defines how EXP, DSCP or Dot1p bits are mapped to internal forwarding class and profile state. The forwarding class and profile state define the Per Hop Behavior (PHB) or the QoS treatment through the 7210 SAS . From release 4.0, the profile mapping is defined using a new policy **mpls-lsp-exp-profile-map**. The **mpls-lsp-exp-profile-map** defines the mapping between the LSP EXP bits and the profile (in or out) to be associated with a packet. The mapping on each **ip-interface** or **port** defaults to the mappings defined in the default network QoS policy until an explicit policy is defined for the IP interface or port. It also defines the rate-limiting parameters for the traffic mapped to each forwarding classes. Traffic mapped to each forwarding class can be rate limited using separate meters for each unicast and multipoint traffic (multipoint is used only for IP Interface for MPLS traffic).

The total number of QoS resources available for network IP interfaces are 512 entries and 256 meters. The software allocates these resources to an IP interface on a first-come first-served basis. The number of IP interfaces that can be successfully configured, varies based on the number of meters and entries utilised per IP interface. Irrespective of the QoS resource allocation, the system enforces the limit of maximum number of IP interfaces to 64.

The total number of QoS resources, that is ingress classification entries and policers, available for use with IP interfaces is limited. The software allocates these resources to an IP interface on a first come first serve basis. The number of resources used per IP interface limits the total number of IP interfaces configured on the system (the total number of IP interfaces allowed is also subject to a system limit).

The egress component of the network QoS policy defines the LSP EXP, DSCP or Dot1p bits marking values associated with each forwarding class.

By default, network qos policy remarking is always disabled. If the egressing packet originated on an ingress SAP, the egress EXP bit marking based on the forwarding class and the profile state. The default map of FC-EXP marking is as shown in default network qos policy, policy id 2. All non-default network qos policies inherits the FC-EXP map.

By default, all ports configured in network mode use Default network policy "1" and all network port IP interfaces use Default network policy "2". Default network policies "1" and "2" cannot be modified or deleted.

Network **policy-id 2** exists as the default policy that is applied to all IP interface by default. The network **policy-id 2** cannot be modified or deleted. It defines the default LSP EXP-to-FC mapping and default meters for unicast and multipoint meters for the ingress MPLS packets. For the egress, it defines eight forwarding classes which defines LSP EXP values and the packet marking criteria.

New (non-default) network policy parameters can be modified. The **no** form of the command reverts the object to the default values.

Changes made to a policy are applied immediately to all IP interface where the policy is applied. For this reason, when a policy requires several changes, it is recommended that you copy the policy to a work area policy-id. The work-in-progress copy can be modified until all the changes are made and then the original policy-id can be overwritten with the **config qos copy** command.

For information about the tasks and commands necessary to access the command line interface and to configure and maintain your devices, refer to CLI Usage chapter in the OS Basic System Configuration Guide.

Normal QoS Operation

The following types of QoS mapping decisions are applicable on a network ingress IP interface .

- MPLS LSP EXP value mapping to FC (if defined)
- Default QoS mapping
- MPLS LSP EXP mapping to profile

The default QoS mapping always exists on an ingress IP interface and every received packet will be mapped to this default if another explicitly defined matching entry does not exist.

Network Qos Policy (ip-interface type) Functionality

When no `ldp-local-fc-enable` is in use, the following restrictions apply (compatible with behavior in release 3.0 and before):

- For LDP LSP traffic, the system uses a single policy (default network policy 2) to assign the FC & profile to the packet. User cannot modify the EXP bits to FC mapping defined in the default policy. Using MPLS EXP bits received in the packet to match the EXP bits configured in the network policy, the policer to use is known. Thus, the system in effect supports only a single system-defined classification policy for all IP interfaces for LDP LSPs though policers with different rates can be used. It does not allow for the flexibility to use different classification policy on different IP interfaces for LDP LSP traffic.
- When using only LDP LSPs, user needs to be aware that the LSP EXP bits to FC classification specified in the global policy is used by the system and not the classification map specified in the network qos policy. Only the meter/policer is used from the network qos policy. It does not cause any issues when only LDP is in use, though the way the policies have been defined currently, its usage is not intuitive to user and can potentially result in confusion.
- There are only 32 unique hardware resources available to map the MPLS EXP bits to profile values. Hence, in release 3.0, only 32 unique network policies could be associated with IP interfaces, though the platform supports more number of IP interfaces. In other words, in release 3.0, multiple IP interfaces needed to share network policies. Though it is sufficient and meets most of the network deployment scenarios, some of the IP interfaces need to share a single network policy. IP interfaces that share the policy will use the same EXP bits to FC mapping.
- If a user receives traffic on RSVP LSP and a LDP LSP with the same value in the EXP bits and on the same IP interface, then potentially, the system can classify packets received on a RSVP LSP to one FC and classify packets received on a LDP LSP to another different FC. This is possible as LDP LSPs use the global network policy for classification and RSVP LSPs used the network policy associated with the IP interface for classification. If both of these policies specify the same EXP to FC mapping then there are no issues. With release 3.0, it is good practice to setup the EXP bits to FC map, to be the same in the default network policy and user-defined network policy when both LDP and RSVP protocols are in use. This is required to ensure that all packets received on an IP interface with a similar value of LSP EXP bit is classified to the same FC and a single policer is used for all them. This ensures that all MPLS packets received with the same EXP bits receive the same QoS treatment in the system.

Executing the command `ldp-use-local-fc-enable`, changes the mode of system behavior for network qos policies, which allows more flexibility and removes some of the restrictions listed above. After executing `ldp-use-local-fc-enable`, the following is possible:

- LSPs setup using LDP uses a global `mpls-lsp-exp-profile-map` policy. By default, the system assigns a default `mpls-lsp-exp-profile-map` policy. User has an option to change

the global policy to use. A new policy `mpls-lsp-exp-profile-map` policy allows the user to assign different profile value for MPLS EXP bits for MPLS packets received over different IP interface. This is helpful for use with primarily RSVP LSP with FRR 1:1. For LDP LSPs or when using FRR facility it is recommended to use a single `mpls-lsp-exp-profile-map` policy for all IP interfaces.

- The new policies separate the profile mapping and FC mapping. The FC used is always picked from the network policy. The new policy syntax allows for flexibility and is intuitive.
- Each IP interface can define a unique network policy for it use, each possibly using a different mapping for MPLS LSP EXP bits to forwarding class (FC). It allows for use of more than 32 distinct network policies, provided network classification resources are available for use.
- If user receives traffic on RSVP LSP and LDP LSP with the same value in the EXP bits, the system provides the same QoS treatment. The system always uses the FC and the meter from the network Qos policy for all MPLS traffic received on an IP interface irrespective of whether its LDP or RSVP LSP.

User cannot execute `'no ldp-use-local-fc-enabled'` after executing `'ldp-use-local-fc-enable'`.

Upgrading from Release 3.0 to Release 4.0

User can use a release 3.0 config file and bootup with release 4.0. The system takes care of the following automatically:

1. System executes the command "no ldp-use-local-fc-enable" and continues to operate with release 3.0 behavior.
2. System changes user defined network policies present in the configuration file to use the new release 4.0 policy commands and parameters. For every network policy, the system creates a new mpls-lsp-exp-profile-map policy with same ID as the network policy ID and associates it with the network policy (using the command mpls-lsp-exp-profile in the network policy ingress context). User cannot modify this association with no ldp-use-local-fc-enable in use.
3. System sets the value of the use-global-mpls-lsp-exp-profile to the default mpls-lsp-exp-profile-map policy "1". The user cannot modify it with no ldp-use-local-fc-enable in use.
4. User is all set to go. They can continue with release 3.0 behavior and continue to use the network policies as before with the restrictions it imposes. It is highly recommended to use the new behavior, as release 3.0 behavior will be deprecated in a future release.
5. User can execute the command 'ldp-use-local-fc-enable' to use the flexibility provided by the new policies. Please take a backup of the existing configuration file before this since once this command is executed and user modifies the existing policies, they cannot use the command 'no ldp-use-local-fc-enable' (release 3.0 behavior) anymore.

DSCP Marking CPU Generated Traffic

DSCP marking for CPU generated traffic is not configurable by the user. The default values are given in [Table 23](#):

Table 23: DSCP and Dot1p Marking

Protocol	IPv4	DSCP Marking	Dot1P Marking	Default FC	DSCP Values	DOT1P Values
OSPF	Yes	Yes	Yes	NC	48	7
ISIS	Yes	Yes	Yes	NC	-	7
TLDP	Yes	Yes	Yes	NC	48	7
RSVP	Yes	Yes	Yes	NC	48	7
SNMP	Yes	Yes	Yes	H2	34	4
NTP	Yes	Yes	Yes	NC	48	7
TELNET	Yes	Yes	Yes	H2	34	4
FTP	Yes	Yes	Yes	H2	34	4
TFTP	Yes	Yes	Yes	H2	34	4
SYSLOG	Yes	Yes	Yes	H2	34	4
TACACS	Yes	Yes	Yes	H2	34	4
RADIUS	Yes	Yes	Yes	H2	34	4
SSH	Yes	Yes	Yes	H2	34	4
ICMP Req	Yes	Yes	Yes	NC	0	7
ICMP Res	Yes	Yes	Yes	NC	0	7
ICMP Unreach	Yes	Yes	Yes	NC	0	7
SCP	Yes	Yes	Yes	H2	34	4
STP	NA	NA	Yes	NC	-	7
CFM	NA	NA	Yes	NC	-	7
ARP	NA	NA	Yes	NC	-	7
Trace route	Yes	Yes	Yes	NC	0	7

Table 23: DSCP and Dot1p Marking (Continued)

Protocol	IPv4	DSCP Marking	Dot1P Marking	Default FC	DSCP Values	DOT1P Values
TACPLUS	Yes	Yes	Yes	H2	34	4
DNS	Yes	Yes	Yes	H2	34	4
BGP	Yes	Yes	Yes	NC	48	7

Note: DSCP and Dot1P values in the table are applicable when remarking is disabled at port level.

Default DSCP Mapping Table

DSCP Name	DSCP Value Decimal	DSCP Value Hexadecimal	DSCP Value Binary	Label
=====				
Default	0	0x00	0b000000	be
nc1	48	0x30	0b110000	h1
nc2	56	0x38	0b111000	nc
ef	46	0x2e	0b101110	ef
af11	10	0x0a	0b001010	assured
af12	12	0x0c	0b001100	assured
af13	14	0x0e	0b001110	assured
af21	18	0x12	0b010010	l1
af22	20	0x14	0b010100	l1
af23	22	0x16	0b010110	l1
af31	26	0x1a	0b011010	l1
af32	28	0x1c	0b011100	l1
af33	30	0x1d	0b011110	l1
af41	34	0x22	0b100010	h2
af42	36	0x24	0b100100	h2
af43	38	0x26	0b100110	h2
default*	0			

*The default forwarding class mapping is used for all DSCP names/values for which there is no explicit forwarding class mapping.

Basic Configurations

A basic network QoS policy must conform to the following:

- Each network QoS policy must have a unique policy ID.
 - Specify the default-action.
 - Have a QoS policy scope of template or exclusive.
 - Have at least one default unicast forwarding class meter.
 - Have at least one multipoint forwarding class meter.
-

Create a Network QoS Policy

Configuring and applying QoS policies other than the default policy is optional. A default network policy of the type 'ip-interface' is applied to each of the **ip-interface** type.

To create an network QoS policy of type ip-interface , define the following:

- A network policy ID value. The system will not dynamically assign a value.
- Set the network-policy-type parameter to be ip-interface.
- Include a description. The description provides a brief overview of policy features.
- You can modify egress LSP EXP marking map. Otherwise, the default values are applied.
 - Remarking — When enabled, this command remarks ALL packets that egress on the specified network port. The remarking is based on the forwarding class to LSP EXP bit mapping defined in the remark policy and associated under the egress node of the network QoS policy.
 - Forwarding class criteria — The forwarding class name represents an egress queue. Specify forwarding class criteria to define the egress characteristics of the queue and the marking criteria of packets flowing through it.
 - LSP EXP — The EXP value is used for all MPLS labeled packets requiring marking that egress on this forwarding class queue that are *in* or *out* of profile.
- Ingress criteria —
 - Default action — Defines the default action to be taken for packets that have an undefined bits set. The default-action specifies the forwarding class to which such packets are assigned.
 - LSP EXP — Creates a mapping between the LSP EXP bits of the network ingress traffic and the forwarding class. Ingress traffic that matches the specified LSP EXP bits will be assigned to the corresponding forwarding class.

User has an option to specify the mapping of the LSP EXP bits to a profile (in/out). Ingress traffic that matches the specified EXP bits will be assigned the corresponding profile.

To create an network QoS policy of type **port**, define the following:

- A network policy ID value. The system will not dynamically assign a value.
- Set the network-policy-type parameter to 'port'
- Include a description. The description provides a brief overview of policy features.
- You can modify egress DSCP and Dot1p marking map. Otherwise, the default values are applied.
 - Remarking — When enabled, this command remarks ALL packets that egress on the specified network port. The remarking is based on the forwarding class to DSCP bit mapping defined in the remark policy and associated under the egress node of the network QoS policy for all IP traffic and forwarding class to Dot1p bit mapping for all IP and MPLS traffic.
 - Forwarding class criteria — The forwarding class name represents an egress queue. Specify forwarding class criteria to define the egress characteristics of the queue and the marking criteria of packets flowing through it.
 - DSCP and Dot1p — The DSCP and Dot1p value is used for all packets requiring marking that egress on this forwarding class queue that are in or out of profile.
- Ingress criteria — Specifies either DSCP or Dot1p (but not both) to forwarding class mapping for all packets.
 - Default action — Defines the default action to be taken for packets that have an undefined DSCP or Dot1p bits set. The default-action specifies the forwarding class to which such packets are assigned.
 - DSCP or Dot1p — Creates a mapping between the DSCP or Dot1p bits of the network ingress traffic and the forwarding class. Ingress traffic that matches the specified DSCP or Dot1p bits will be assigned to the corresponding forwarding class.

Use the following CLI syntax to create a network QoS policy:

CLI Syntax: `config>qos#`

```

network policy-id [network-policy-type network-policy-type]
description description-string
scope {exclusive|template}
egress
    remarking
    remark <policy-id>
ingress
    default-action fc {fc-name} profile {in|out}
    lsp-exp lsp-exp-value fc fc-name profile {in | out}
    fc {fc-name}
    meter {meter-id}
```

Basic Configurations

```
multicast-meter {id}
meter meter-id [multipoint]
adaptation-rule cir {closest | max | min} pir {closest | max | min}
cbs {size-in-kbits}
mbs {size-in-kbits}
mode {trtcm | srtcm}
rate cir cir-rate-in-kbps [pir pir-rate-in-kbps]
mpls-lsp-exp-profile policy-id
```

```
config>qos>network# info
-----
description "Network Qos policy 200"
ingress
    meter 1 create
    exit
    meter 9 multipoint create
    exit
exit
egress
    remarking
exit
-----
A:ALA-10config>qos>network#
```

The following output displays the configuration for router interface ALA-1-2 with network policy 600 applied to the interface.

```
A:ALA-7>config>router# info
#-----
echo "IP Configuration"
#-----
...
    interface "ALA-1-2"
        address 10.10.4.3/24
        qos 600
    exit
...
-----
A:ALA-7>config>router#
```

Default Network Policy Values

The default network policy for IP interfaces is identified as policy-id 2. Default policies cannot be modified or deleted. The following displays default network policy parameters:

Table 24: Network Policy Defaults

Field	Default
description	Default network QoS policy.
scope	template
ingress	
default-action	fc be profile out (default action profile out is applicable only for port policies and not for ip-interface policies)
egress	
remarking	
fc af:	
lsp-exp-in-profile	3
lsp-exp-out-profile	2
fc be:	
lsp-exp-in-profile	0
lsp-exp-out-profile	0
fc ef:	
lsp-exp-in-profile	5
lsp-exp-out-profile	5
fc h1:	
lsp-exp-in-profile	6
lsp-exp-out-profile	6
fc h2:	
lsp-exp-in-profile	4

Table 24: Network Policy Defaults (Continued)

Field	Default
lsp-exp-out-profile	4
fc l1:	
lsp-exp-in-profile	3
lsp-exp-out-profile	2
fc l2:	
lsp-exp-in-profile	1
lsp-exp-out-profile	1
fc nc:	
lsp-exp-in-profile	7
lsp-exp-out-profile	7

Table 25: Default Network QoS Policy of Type IP Interface, LSP EXP to FC Mapping on Ingress (Color aware policing is supported on network ingress.)

LSP EXP Value	7210 FC Ingress	Profile
0	be	Out
1	l2	In
2	af	Out
3	af	In
4	h2	In
5	ef	In
6	h1	In
7	nc	In

The default network policy for port is identified as policy-id 1. Default policies cannot be modified or deleted. The following output displays the parameters for default network policy of type **port** :

```
*A:ALA>config>qos>network# info detail
```

```
-----
```

```
description "Default network-port QoS policy."
scope template
ingress
```

```

default-action fc be profile out
meter 1 create
    mode trtcm
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    mbs default
    cbs default
exit
dscp be fc be profile out
dscp ef fc ef profile in
dscp cs1 fc l2 profile in
dscp nc1 fc h1 profile in
dscp nc2 fc nc profile in
dscp af11 fc af profile in
dscp af12 fc af profile out
dscp af41 fc h2 profile in
exit
egress
    no remarking
    remark 1
exit
-----
*A:ALA>config>qos>network#

```

Default remark policy used for Dot1p and DSCP marking is as shown below:

```

fc af
    dscp-in-profile af11
    dscp-out-profile af12
    no lsp-exp-in-profile
    no lsp-exp-out-profile
    no dot1p-lsp-exp-in-profile
    no dot1p-lsp-exp-out-profile
    dot1p-in-profile 3
    dot1p-out-profile 2
exit
fc be
    dscp-in-profile be
    dscp-out-profile be
    no lsp-exp-in-profile
    no lsp-exp-out-profile
    no dot1p-lsp-exp-in-profile
    no dot1p-lsp-exp-out-profile
    dot1p-in-profile 0
    dot1p-out-profile 0
exit
fc ef
    dscp-in-profile ef
    dscp-out-profile ef
    no lsp-exp-in-profile
    no lsp-exp-out-profile
    no dot1p-lsp-exp-in-profile
    no dot1p-lsp-exp-out-profile
    dot1p-in-profile 5
    dot1p-out-profile 5
exit
fc h1
    dscp-in-profile nc1

```

```
dscp-out-profile nc1
no lsp-exp-in-profile
no lsp-exp-out-profile
no dot1p-lsp-exp-in-profile
no dot1p-lsp-exp-out-profile
dot1p-in-profile 6
dot1p-out-profile 6
exit
fc h2
dscp-in-profile af41
dscp-out-profile af41
no lsp-exp-in-profile
no lsp-exp-out-profile
no dot1p-lsp-exp-in-profile
no dot1p-lsp-exp-out-profile
dot1p-in-profile 4
dot1p-out-profile 4
exit
fc l1
dscp-in-profile af21
dscp-out-profile af22
no lsp-exp-in-profile
no lsp-exp-out-profile
no dot1p-lsp-exp-in-profile
no dot1p-lsp-exp-out-profile
dot1p-in-profile 3
dot1p-out-profile 2
exit
fc l2
dscp-in-profile cs1
dscp-out-profile cs1
no lsp-exp-in-profile
no lsp-exp-out-profile
no dot1p-lsp-exp-in-profile
no dot1p-lsp-exp-out-profile
dot1p-in-profile 1
dot1p-out-profile 1
exit
fc nc
dscp-in-profile nc2
dscp-out-profile nc2
no lsp-exp-in-profile
no lsp-exp-out-profile
no dot1p-lsp-exp-in-profile
no dot1p-lsp-exp-out-profile
dot1p-in-profile 7
dot1p-out-profile 7
exit
```

Service Management Tasks

Deleting QoS Policies

A network policy is associated by default with IP interfaces and network ports.

You can replace the default policy with a non-default policy, but you cannot remove default policies from the configuration. When you remove a non-default policy, the policy association reverts to the appropriate default network policy.

Remove a Policy from the QoS Configuration

To delete a network policy, enter the following commands:

CLI Syntax: `config>qos# no network network-policy-id`

Copying and Overwriting Network Policies

You can copy an existing network policy to a new policy ID value or overwrite an existing policy ID. The `overwrite` option must be specified or an error occurs if the destination policy ID exists.

Note: In "no ldp-local-fc-enable" mode, when an ip-interface qos policy is copied, a new mpls-lsp-exp-profile-map is created with the id same as dest-policy id and the mpls-lsp-exp-profile of the dest-policy points to it. In "ldp-local-fc-enable" mode, when an ip-interface qos policy is copied, the mpls-lsp-exp-profile of the dest-policy points to the mpls-lsp-exp-profile-map pointed to by the source-policy.

CLI Syntax: `config>qos# copy network source-policy-id dest-policy-id [overwrite]`

The following output displays the copied policies:

```
A:ALA-12>config>qos# info detail
-----
...
    network 1 create
      description "Default network QoS policy."
      scope template
      ingress
      default-action fc be profile out
...
    network 600 create
      description "Default network QoS policy."
      scope template
      ingress
      default-action fc be profile out
...
    network 700 create
      description "Default network QoS policy."
      scope template
      ingress
      default-action fc be profile out
...
-----
A:ALA-12>config>qos#
```

Editing QoS Policies

You can change existing policies, except the default policies, and entries in the CLI. The changes are applied immediately to all network where the policy is applied. To prevent configuration errors use the copy command to make a duplicate of the original policy to a work area, make the edits, and then overwrite the original policy. The number of meters (TP) used are: 5 (Meters 1,2,3,9,12).

Resource Allocation for Network QoS policy

This section describes the allocation of QoS resources for network QoS policy (for type=ip interface only).

When an IP interface is created, a default network QoS policy is applied. For the default policy, two meters and two classification entries in hardware are allocated.

The resources are allocated to a network policy, only when a port is configured for the IP interface.

For every FC in use, the system allocates two classification entries in hardware. If multiple matchcriteria entries map to the same FC, then each of these are allocated two classification entries in hardware. For example, if there are two match-criteria entries that map to FC 'af', then a total of four classification entries are allocated in hardware and if there are four match-criteria entries that map to FC 'af', then a total of 8 classification entries are allocated in hardware.

For every meter or policer in use, the system allocates one meter in hardware. A meter or policer is considered to be in use when it is associated with an FC in use.

The number of IP interfaces allowed is limited to number of resources available in hardware, subject to system limit (a maximum of 64 IP interfaces are allowed). The system reserves a total of 512 classification entries and 256 meters in hardware for use by network policy associated with an IP interface.

For computing the number of QoS resources used by an IP interface:

- Determine number of match-criteria entries used to identify the FC.
- Determine number of FCs to use.

Only the FCs used by the match-criteria classification entries are to be considered for the 'number of FCs'. Therefore are referred to as 'FC in use'.

Use the following rules to compute the number of classification entries per FC in use:

If a FC is in use and is created without explicit meters, use default meter#1 for unicast traffic and default meter #9 for all other traffic types (that is, broadcast, multicast and unknown-unicast). This requires two classification entries in hardware.

If a FC is in use and is created with an explicit unicast meter, use that meter for unicast traffic and use default meter #9 for all other traffic types. This requires two classification entries in hardware.

If a FC is in use and is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for all other kinds of traffic. This requires two classification entries in hardware.

Given the number of match criteria and the number of FCs used, use the equation given below to arrive at total number of classification entries per policy (for example TC):

$$TC = \sum_{i=nc,h1,ef,h2,l1,af,l2,be} 2 * E(i)$$

Where,

E(i) is the number of match- criteria entries that classify packets to FCi. For 7210 platforms, the maximum number of classification entries per policy can be 64 (including default).

2 is the number of classification entries that are required by FCi.

Note: In any case, only 2 classification entries are used per FC in a network policy, as only two traffic-types are supported.

Determine number of policers or meters to use (for example TP). A maximum of 12 meters per network policy is available.

Only those meters that are associated with FCs need to be considered for number of meters. Note, that only FCs in use are considered.

Network QoS Policies Resource Usage Examples

Example 1

```
network 1 network-policy-type ip-interface create
  description "network-policy-1"
    ingress
      default-action fc be
      meter 1 create
      exit
      meter 9 multipoint create
      exit
    exit
  egress
    exit
```

The number of classification entries (TC) used is calculated, as follows:

$$(2 * 0)nc + (2 * 0)h1 + (2 * 0)ef + (2 * 0)h2 + (2 * 0)l1 + (2 * 0)af + (2 * 0)l2 + (2 * 1)be = 2$$

The number of meters (TP) used are: 2 (meter 1 and 9).

Example 2

```
network 2 network-policy-type ip-interface create
  description "network-policy-2"
    ingress
      default-action fc be
      meter 1 create
      exit
      meter 2 create
      exit
      meter 9 multipoint create
      exit
      meter 12 multipoint create
      exit
      fc "af" create
        meter 2
        multicast-meter 12
      exit
      lsp-exp 2 fc af
    exit
  egress
    exit
```

The number of classification entries (TC) used is calculated, as follows:

$$(2 * 0)nc + (2 * 0)h1 + (2 * 0)ef + (2 * 0)h2 + (2 * 0)l1 + (2 * 1)af + (2 * 0)l2 + (2 * 1)be = 4$$

The number of meters (TP) user are: 4 (Meters 1,2,9,12)

Example 3

```

network 3 network-policy-type ip-interface create
  description "network-policy-3"
  ingress
    default-action fc be
    meter 1 create
    exit
    meter 2 create
    exit
    meter 9 multipoint create
    exit
    meter 12 multipoint create
    exit
    fc "af" create
      meter 2
      multicast-meter 12
    exit
    fc "be" create
      meter 2
      multicast-meter 12
    exit
    lsp-exp 2 fc af
  exit
egress
exit

```

The number of classification entries (TC) used are calculated, as follows:

$$(2 * 0)nc + (2 * 0)h1 + (2 * 0)ef + (2 * 0)h2 + (2 * 0)l1 + (2 * 1)af + (2 * 0)l2 + (2 * 1)be = 4$$

The number of meters (TP) user are: 2 (Meters 2,12).

Example 4

```
network 4 network-policy-type ip-interface create
  description "network-policy-4"
  ingress
    default-action fc be
    meter 1 create
    exit
    meter 9 multipoint create
    exit
    lsp-exp 1 fc l2
    lsp-exp 2 fc af
    lsp-exp 3 fc af
    lsp-exp 4 fc h2
    lsp-exp 5 fc ef
    lsp-exp 6 fc h1
    lsp-exp 7 fc nc
  exit
  egress
  exit
exit
```

The number of Filter-Entries (TC) used is calculated, as follows:

$$(2 * 1)nc + (2 * 1)h1 + (2 * 1)ef + (2 * 1)h2 + (2 * 0)l1 + (2 * 2)af + (2 * 1)l2 + (2 * 1)be = 16$$

The number of meters (TP) used are: 2 (Meters 1,9).

Example 5

```
network 5 network-policy-type ip-interface create
  description "network-policy-5"
  ingress
    default-action fc be
    meter 1 create
    exit
    meter 2 create
    exit
    meter 9 multipoint create
    exit
    meter 12 multipoint create
    exit
    fc "af" create
    exit
    fc "be" create
    exit
    fc "ef" create
    exit
    fc "h1" create
    exit
    fc "h2" create
```

```

exit
fc "l2" create
exit
fc "nc" create
exit
lsp-exp 1 fc l2
lsp-exp 2 fc af
lsp-exp 3 fc af
lsp-exp 4 fc h2
lsp-exp 5 fc ef
lsp-exp 6 fc h1
lsp-exp 7 fc nc
exit
egress
exit

```

The number of classification entries (TC) used is calculated, as follows:

$$(2 * 1)nc + (2 * 1)h1 + (2 * 1)ef + (2 * 1)h2 + (2 * 0)l1 + (2 * 2)af + (2 * 1)l2 + (2 * 1)be = 16$$

The number of meters (TP) used are: 2 (Meters 1,9 – Note that meters 2 and 12 are not accounted for, since its not associated with any FC).

Example 6

```
network 6 network-policy-type ip-interface create
description "network-policy-6"

    ingress
        default-action fc be
        meter 1 create
        exit
        meter 2 create
        exit
        meter 3 create
        exit
        meter 9 multipoint create
        exit
        meter 12 multipoint create
        exit
        fc "af" create
            meter 2
            multicast-meter 12
        exit
        fc "be" create
        exit
        fc "ef" create
        exit
        fc "h1" create
            meter 3
        exit
        fc "h2" create
        exit
        fc "l2" create
        exit
        fc "nc" create
            meter 3
        exit
        lsp-exp 1 fc l2
        lsp-exp 2 fc af
        lsp-exp 3 fc af
        lsp-exp 4 fc h2
        lsp-exp 5 fc ef
        lsp-exp 6 fc h1
        lsp-exp 7 fc nc
    exit
    egress
    exit
exit
```

The number of classification entries (TC) used is calculated, as follows:

$$(2 * 1)nc + (2 * 1)h1 + (2 * 1)ef + (2 * 1)h2 + (2 * 0)l1 + (2 * 2)af + (2 * 1)l2 + (2 * 1)be = 16$$

The number of meters (TP) used are: 5 (Meters 1,2,3,9,12).

Example 7

```

network 2 network-policy-type ip-interface create
  description "Default network QoS policy."
  scope template
  ingress
    default-action fc be
    meter 1 create
      mode trtcm
      adaptation-rule cir closest pir closest
      rate cir 0 pir max
      mbs default
      cbs default
    exit
    meter 9 multipoint create
      mode trtcm
      adaptation-rule cir closest pir closest
      rate cir 0 pir max
      mbs default
      cbs default
    exit
    lsp-exp 0 fc be
    lsp-exp 1 fc l2
    lsp-exp 2 fc af
    lsp-exp 3 fc af
    lsp-exp 4 fc h2
    lsp-exp 5 fc ef
    lsp-exp 6 fc h1
    lsp-exp 7 fc nc
  exit
  egress
    no remarking
  exit
exit

```

The number of classification entries (TC) used is: 2.

The number of meters (TP) used is: 2.

Example 8

```

network 8 network-policy-type ip-interface create
  description "network-policy-8"
  ingress
    default-action fc nc
    meter 1 create
    exit
    meter 2 create
    exit
    meter 3 create
    exit
    meter 4 create
    exit
    meter 5 create
    exit
    meter 7 multipoint create

```

Service Management Tasks

```
exit
meter 8 multipoint create
exit
meter 9 multipoint create
exit
meter 12 multipoint create
exit
fc "af" create
    meter 2
    multicast-meter 12
exit
fc "ef" create
    meter 4
    multicast-meter 8
exit
fc "h2" create
exit
fc "l2" create
    meter 3
    multicast-meter 7
exit
fc "nc" create
    meter 4
    multicast-meter 8
exit
lsp-exp 1 fc l2
lsp-exp 3 fc af
lsp-exp 5 fc ef
lsp-exp 7 fc nc
exit
egress
exit
exit
```

The number of classification entries (TC) used is calculated, as follows:

$$(2 * 2)_{nc} + (2 * 0)_{h1} + (2 * 1)_{ef} + (2 * 0)_{h2} + (2 * 0)_{l1} + (2 * 1)_{af} + (2 * 1)_{l2} + (0 * 0)_{be} = 10$$

The numbers of meters (TP) used is: 6 (Meters 2, 3, 4, 7, 8, 12).

Network QoS Policy Command Reference

Command Hierarchies

- [Configuration Commands on page 115](#)
- [Operational Commands on page 116](#)
- [Show Commands on page 116](#)

Configuration Commands

```
config
  — qos
    — [no] mpls-lsp-exp-profile-map policy-id [create]
      — description description-string
      — no description
      — lsp-exp lsp-exp-value profile {in|out}
      — no lsp-exp
    — [no] ldp_local_fc_enable
    — [no] use-global-mpls-lsp-exp-profile policy-id

config
  — qos
    — [no] network network-policy-id [network-policy-type { ip-interface | port } ]
      — description description-string
      — no description
      — scope {exclusive | template}
      — no scope
      — egress
        — [no] remark <policy-id>
        — [no] remarking
      — ingress
        — default-action fc fc-name profile {in | out}
        — dot1p dot1p-priority fc fc-name profile {in | out | use-de}
        — no dot1p dot1p-priority
        — [no] fc fc-name [create]
          — meter meter-id
          — no meter
          — multicast-meter meter-id
          — no multicast-meter
        — dscp dscp-name fc fc-name profile {in | out}
        — no dscp dscp-name
        — lsp-exp lsp-exp-value fc fc-name profile {in | out}
        — no lsp-exp lsp-exp-value
        — meter meter-id [multipoint] [create]
        — no meter meter-id
```

- **adaptation-rule** [**cir** *adaptation-rule*] [**pir** *adaptation-rule*]
- **no adaptation-rule**
- **cbs** *size-in-kbits*
- **no cbs**
- **mbs** *size-in-kbits*
- **no mbs**
- **mode** {**trtcn** 1| **srtcm**}
- **no mode**
- **rate** *cir-rate-in-kbps* [**pir** *pir-rate-in-kbps*]
- **no rate**
- **[no] mpls-lsp-exp-profile-map** *policy-id*

Operational Commands

- config**
 - **qos**
 - **copy network** *src-pol dst-pol* [**overwrite**]

Show Commands

- show**
 - **qos**
 - **network** *policy-id* [**detail**]
 - **mpls-lsp-exp-profile-map** [*policy-id*] [**detail**]

Configuration Commands

Generic Commands

description

Syntax **description** *description-string*
 no description

Context config>qos>network *policy-id*
 config>qos>mpls-lsp-exp-profile-map

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Default No description is associated with the configuration context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands

copy

Syntax	copy network <i>src-pol dst-pol</i> [overwrite]
Context	config>qos
Description	This command copies existing QoS policy entries for a QoS policy-id to another QoS policy-id. The copy command is used to create new policies using existing policies and also allows bulk modifications to an existing policy with the use of the overwrite keyword. Note: In "no ldp-local-fc-enable" mode, when an ip-interface qos policy is copied, a new mpls-lsp-exp-profile-map is created with the id same as dest-policy id and the mpls-lsp-exp-profile of the dest-policy points to it. In "ldp-local-fc-enable" mode, when an ip-interface qos policy is copied, the mpls-lsp-exp-profile of the dest-policy points to the mpls-lsp-exp-profile-map pointed to by the source-policy.
Parameters	network <i>src-pol dst-pol</i> — Indicates that the source and destination policies are network policy IDs. Specify the source policy that the copy command will copy and specify the destination policy to which the command will duplicate the policy to a new or different policy ID. Values 1 — 65535 overwrite — Specifies to replace the existing destination policy. Everything in the existing destination policy will be overwritten with the contents of the source policy. If overwrite is not specified, an error will occur if the destination policy ID exists. <pre>SR>config>qos# copy network 1 427 MINOR: CLI Destination "427" exists use {overwrite}. SR>config>qos# copy network 1 427 overwrite</pre>

scope

Syntax	scope { exclusive template } no scope
Context	config>qos>network <i>policy-id</i>
Description	This command configures the network policy scope as exclusive or template. The no form of this command sets the scope of the policy to the default of template.
Default	template

- Parameters**
- exclusive** — When the scope of a policy is defined as exclusive, the policy can only be applied to one interface. If a policy with an exclusive scope is assigned to a second interface an error message is generated. If the policy is removed from the exclusive interface, it will become available for assignment to another exclusive interface.
The system default policies cannot be put into the exclusive scope. An error will be generated if scope exclusive is executed in any policies with a policy-id equal to 1.
 - template** — When the scope of a policy is defined as template, the policy can be applied to multiple interfaces on the router.
Default QoS policies are configured with template scopes. An error is generated if you try to modify the template scope parameter to exclusive scope on default policies.

Network QoS Policy Commands

network

Syntax **[no]** **network** *network-policy-id* [**network-policy-type** { **ip-interface** | **port** }]

Context config>qos

Description This command creates or edits a QoS network policy. The network policy defines the treatment packets receive as they ingress and egress the network port.

The QoS network policy consists of an ingress and egress component. The ingress component of the policy defines how bits are mapped to internal forwarding class and profile state. The forwarding class and profile state define the Per Hop Behavior (PHB) or the QoS treatment through the 7210 SAS. The mapping on each network interface defaults to the mappings defined in the default network QoS policy until an explicit policy is defined for the network interface. It also defines the rate-limiting parameters for the traffic mapped to each forwarding classes. Traffic mapped to each forwarding class can be rate limited using separate meters for each uni-cast and multipoint traffic.

The egress component of the network QoS policy defines forwarding class and profile state to LSP EXP values for traffic to be transmitted into the core network. If the egressing packet originated on an ingress SAP, the parameter is for the network, the egress QoS policy also defines the bit marking based on the forwarding class and the profile state.

Network **policy-id 2** exists as the default policy that is applied to all network interfacesIP interface by default. The network **policy-id 2** cannot be modified or deleted. It defines the default LSP EXP-to-FC mapping and default meters for unicast and multipoint metersfor the ingressMPLS packets. For the egress, it defines eight forwarding classes which defines LSP EXP values and the packet marking criteria.

Network policy-id 1 exists as the default policy that is applied to all network ports by default. This default policy cannot be modified or deleted. It defined the default DSCP-to-FC mapping and default unicast meters for ingress IP traffic. For the egress, if defines the forwarding class to Dot1p and DSCP values and the packet marking criteria.

If a new network policy is created (for instance, policy-id), only the default actiondefault meters for unicast and multipoint traffic and egress forwarding class parameters are identical to the default policy. A new network policy does not contain the default LSP EXP-to-FC mapping for network QoS policyof type **ip-interface** or the DSCP-to-FC mapping (for network QoS policy of type **port**). The default network policy can be copied (use the copy command) to create a new network policy that includes the default ingress LSP EXP or DSCP to FC mapping (as appropriate). You can modify parameters or use the **no** modifier to remove an object from the configuration.

Any changes made to an existing policy, using any of the sub-commands, will be applied immediately to all network ports where this policy is applied. For this reason, when many changes

are required on a policy, it is highly recommended that the policy be copied to a work area policy-id. That work-in-progress policy can be modified until complete and then written over the original policy-id. Use the `config qos copy` command to maintain policies in this manner.

The **no** form of this command deletes the network policy. A policy cannot be deleted until it is removed from all entities where it is applied. The default network **policy** *policy-id* 1 cannot be deleted.

Default System Default Network Policy 1

Parameters *network-policy-id* — The policy-id uniquely identifies the policy on the .

Default none

Values 1— 65535

network-policy-type — The type of the policy, either **ip-interface** or **port**. It defines where this network policy can be applied.

ip-interface — Specifies only EXP-based classification rules and marking values. It can only be associated with an IP interface.

port — Specifies only DSCP and Dot1p classification rules and marking values. It can only be associated with a portde.

mpls-lsp-exp-profile-map

Syntax **mpls-lsp-exp-profile-map policy-id [create]**
no mpls-lsp-exp-profile-map

Context config>qos

Description This command allows the user to create a new mpls-lsp-exp-profile-map policy. The policy specifies the profile to assign to the packet based on the MPLS LSP EXP bits value matched in the MPLS packet received on a network IP interface.

The assigned profile is available for use by the meter/policer associated with FC in the network policy attached to this IP interface.

The policy is associated with network policy attached to a network IP interface.

When 'no ldp-use-local-fc-enable' is set, system creates the mpls-lsp-exp-profile-map automatically with same ID as the network policy ID. The values that map the lsp-exp bits to a profile value can be modified by the user. The system deletes the policy when the associated network policy is deleted.

When ldp-use-local-fc-enable is set, system does not create the mpls-lsp-exp-profile-map policies by default (except for the default policy "1"). User is allowed to create, delete, modify, copy the

policies. User needs to associate these policies with appropriate network policies as per their requirement.

Default 1 (default mpls-lsp-exp-profile-map policy “1”).

Parameters *policy-id* — The policy-id uniquely identifies the policy on the 7210 SAS.

Values 1 — 65535

create — The keyword used to create a policy.

lsp-exp

Syntax **lsp-exp** *lsp-exp-value* **profile** {in|out}
no lsp-exp

Context config>qos> mpls-lsp-exp-profile-map

Description This command creates a mapping between the LSP EXP bits of the network ingress traffic and the profile.

Ingress traffic that matches the specified LSP EXP bits will be assigned the corresponding profile.

Multiple commands can be entered to define the association of some or all eight LSP EXP bit values to the profile. For undefined values, packets are assigned the profile value out.

The no form of this command removes the association of the LSP EXP bit value to the profile value. The default profile value ‘out’ then applies to that LSP EXP bit pattern.

Default none

Parameters *lsp-exp-value* — The 3-bit LSP EXP bit value, expressed as a decimal integer.

Values 0 — 7

profile — Assign the profile value to be assigned to this LSP EXP value.

Default None, the lsp-exp command must define a profile state.

Values Values in, out

use-global-mpls-lsp-exp-profile

Syntax **use-global-mpls-lsp-exp-profile** *policy-id*
no use-global-mpls-lsp-exp-profile

Context config>qos

Description This command allows the user to associate the mpls-lsp-exp-profile-map policy for use with LDP LSPs. When color aware metering is in use for the IP interface, the policy specified here provides

the profile to assign to the MPLS packets received on any of the network IP interface in use in the system. The MPLS EXP bits in the received packet are matched for assigning the profile.

When ‘no ldp-use-local-fc-enable’ is set, system sets it to the default value. User cannot modify it.

When ldp-use-local-fc-enable is set, on system boot-up sets it to the default value. User can modify it to use the policy of their choice.

For LDP LSP traffic, the system always uses the global mpls-lsp-exp-profile-map policy. For RSVP LSP traffic, system uses the mpls-lsp-exp-profile-map policy associated with the network policy. It is highly recommended to use a single mpls-lsp-exp-profile-map policy for all the network policies when FRR facility is in use for consistent QoS treatment.

The **no** form of the command sets the policy to default policy.

Default Default mpls-lsp-exp-profile-map policy “1” is used.

Parameters *policy-id* — The policy-id uniquely identifies the mpls-lsp-exp-profile-map policy to use.

Values 1 — 65535

mpls-lsp-exp-profile

Syntax **mpls-lsp-exp-profile policy-id [create]**
no mpls-lsp-exp-profile

Context config>qos>network>ingress

Description Specify the mpls-lsp-exp-profile-map policy to use for assigning profile values for packets received on this IP interface.

When ‘no ldp-use-local-fc-enable’ is set, this policy is managed by the system. User is not allowed to modify it. The system assigns the same policy ID as the network policy ID. It is cannot be modified by the user.

When ‘ldp-use-local-fc-enable’ is set, by default the system assigns the default policy ID “1”. User can create new policies and specify the new policy instead of the default policy.

Note: For LDP LSP traffic, the system always uses the global mpls-lsp-exp-profile-map policy. For RSVP LSP traffic, system uses the mpls-lsp-exp-profile-map policy associated with the network policy. It is highly recommended to use a single mpls-lsp-exp-profile-map policy for all the network policies when FRR facility is in use for consistent QoS treatment.

The **no** form of the command assigns the default policy.

Parameters *policy-id* — The policy-id uniquely identifies the policy on the 7210 SAS.

Values 1 — 65535

ldp_local_fc_enable

Syntax **ldp-local-fc-enable**
 no ldp-local-fc-enable

Context config>qos

Description This command determines the system QoS behavior for network IP interfaces for MPLS traffic.

The **no** form of the command allows for backward compatibility with prior releases. With the **no** form, the system continues to operate with release 3.0 (or before) network QoS behavior. The following restrictions apply when operating with **no** form of the command:

- For LDP LSP traffic, the system uses a single policy (default network policy 1) to assign the FC & profile to the packet. User cannot modify the EXP bits to FC mapping defined in the default policy. Using MPLS EXP bits received in the packet to match the EXP bits configured in the network policy, the policer to use is known. Thus, the system in effect supports only a single system-defined classification policy for all IP interfaces for LDP LSPs though policers with different rates can be used. It does not allow for the flexibility to use different classification policy on different IP interfaces for LDP LSP traffic.
- When using only LDP LSPs, user needs to be aware that the LSP EXP bits to FC classification specified in the global policy is used by the system and not the classification map specified in the network qos policy. Only the meter/policer is used from the network qos policy. It does not cause any issues when only LDP is in use, though the way the policies have been defined currently, its usage is not intuitive to user and can potentially result in confusion.
- There are only 32 unique hardware resources available to map the MPLS EXP bits to profile values. Hence, in release 3.0, only 32 unique network policies could be associated with IP interfaces, though the platform supports more number of IP interfaces. In other words, in release 3.0, multiple IP interfaces needed to share network policies. Though it is sufficient and meets most of the network deployment scenarios, some of the IP interfaces need to share a single network policy. IP interfaces that share the policy will use the same EXP bits to FC mapping.
- If a user receives traffic on RSVP LSP and a LDP LSP with the same value in the EXP bits and on the same IP interface, then potentially, the system can classify packets received on a RSVP LSP to one FC and classify packets received on a LDP LSP to another different FC. This is possible as LDP LSPs use the global network policy for classification and RSVP LSPs used the network policy associated with the IP interface for classification. If both of these policies specify the same EXP to FC mapping then there are no issues. With release 3.0, it is good practice to setup the EXP bits to FC map, to be the same in the default network policy and user-defined network policy when both LDP and RSVP protocols are in use. This is required to ensure that all packets received on an IP interface with a similar value of LSP EXP bit is classified to the same FC and a single policer is used for all them. This ensures that all MPLS packets received with the same EXP bits receive the same QoS treatment in the system.

Executing the command `ldp-use-local-fc-enable`, changes the mode of system behavior which allows more flexibility to the user and removes some of the restrictions listed above. With `ldp-use-local-fc-enable` set, the following is possible:

- LSPs setup using LDP uses a global `mpls-lsp-exp-profile-map` policy. By default, the system assigns a default `mpls-lsp-exp-profile-map` policy. User has an option to change the global policy to use. A new policy `mpls-lsp-exp-profile-map` policy allows the user to assign different profile value for MPLS EXP bits for MPLS packets received over different IP interface. This is helpful for use with primarily RSVP LSP with FRR 1:1. For LDP LSPs or when using FRR facility it is recommended to use a single `mpls-lsp-exp-profile-map` policy for all IP interfaces.
- The new policies separate the profile mapping and FC mapping. The FC used is always picked from the network policy. The new policy syntax allows for flexibility and is intuitive.
- Each IP interface can define a unique network policy for it use, each possibly using a different mapping for MPLS LSP EXP bits to forwarding class (FC). It allows for use of more than 32 distinct network policies, provided network classification resources are available for use.
- If user receives traffic on RSVP LSP and LDP LSP with the same value in the EXP bits, the system provides the same QoS treatment. The system always uses the FC and the meter from the network Qos policy for all MPLS traffic received on an IP interface irrespective of whether its LDP or RSVP LSP.

Note: User cannot execute the command ‘`no ldp-use-local-fc-enable`’ after executing the command ‘`ldp-use-local-fc-enable`’.

Default `no ldp_fc_local_enable`

Network Ingress QoS Policy Commands

ingress

Syntax	ingress
Context	config>qos>network <i>policy-id</i>
Description	This command is used to enter the CLI node that creates or edits policy entries that specify the to forwarding class mapping packets. When pre-marked packets ingress on a network port, the QoS treatment through the 7210 SAS-based on the mapping defined under the current node.

default-action

Syntax	default-action fc <i>fc-name</i> [profile { in out }]
Context	config>qos>network>ingress
Description	This command defines or edits the default action to be taken for packets that have an undefined LSP EXP bits set. The default-action command specifies the forwarding class to which such packets are assigned. Multiple default-action commands will overwrite each previous default-action command.
Default	default-action fc be profile out
Parameters	fc <i>fc-name</i> — Specify the forwarding class name. All packets with LSP EXP or dot1p bits bits that is not defined will be placed in this forwarding class. Default None, the fc name must be specified Values be, l2, af, l1, h2, ef, h1, nc profile {in out} — All packets that are assigned to this forwarding class will be considered in or out of profile based on this command. , on network ingress, the meter/policer supports color-aware policing/ metering. The value of the profile parameter is used to provide the color to the meter. Value of 'in' indicates 'Green' color OR in-profile packet to the meter and value of 'out' indicates 'Yellow' color OR out-of-profile packet to the meter operating in color-aware mode. Based on the configured meter rates, the final profile for the packet is determined. The final color is used for subsequent processing of the packet in the system. On egress, in case of congestion, the in-profile packets are preferentially queued over the out-of-profile packets. The profile can be specified in 3.0 release. Default None Values in, out

dot1p

Syntax	dot1p <i>dot1p-priority</i> fc <i>fc-name</i> profile { in out } no dot1p <i>dot1p-priority</i>
Context	config>qos>network>ingress
Description	This command explicitly sets the forwarding class or enqueueing priority and profile of the packet when a packet is marked with a <i>dot1p-priority</i> specified. Adding a dot1p rule on the policy forces packets that match the <i>dot1p-priority</i> specified to override be assigned to the forwarding class and enqueueing priority and profile of the packet based on the parameters included in the Dot1p rule. The <i>dot1p-priority</i> is derived from the most significant three bits in the IEEE 802.1Q or IEEE 802.1P header. The three dot1p bits define 8 Class-of-Service (CoS) values commonly used to map packets to per-hop Quality-of-Service (QoS) behavior. The no form of this command removes the explicit dot1p classification rule from the policy. Removing the rule on the policy immediately removes the rule on all ingress SAPs ports using the policy.
Parameters	<i>dot1p-priority</i> — This value is a required parameter that specifies the unique IEEE 802.1P value that will match the dot1p rule. If the command is executed multiple times with the same <i>dot1p-value</i>, the previous forwarding class is completely overridden by the new parameters . A maximum of eight dot1p rules are allowed on a single policy. Values 0 — 7 fc <i>fc-name</i> — The value given for the <i>fc-name</i> parameter must be one of the predefined forwarding classes in the system. Specifying the <i>fc-name</i> is optional. When a packet matches the rule, the forwarding class is only overridden when the fc <i>fc-name</i> parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches. Values be, l2, af, l1, h2, ef, h1, nc profile {in out } — All packets that are assigned to this forwarding class will be considered in or out of profile based on this command or to use the default. In case of congestion, the in-profile packets are preferentially queued over the out-of-profile packets. Default none, the profile name must be specified.

meter

Syntax	meter <i>meter-id</i> no meter <i>meter-id</i> [multipoint] [create]
Context	config>qos>network>ingress
Description	This command enables the context to configure an ingress Network QoS policy meter. The meter command allows the creation of multipoint meters. Only multipoint meters can receive ingress packets that need to be sent to multiple destinations. Multipoint meters are for traffic bound to multiple destinations. Within non-multipoint services, such as Epipe services, all traffic is considered unicast due to the nature of the service type. Multicast and broadcast-destined traffic in an Epipe service will not be mapped to a multipoint service meter. The no form of this command removes the meter-id from the Network ingress QoS policy and from any existing Ports using the policy. If any forwarding class forwarding types are mapped to the meter, they revert to their default meters. When a meter is removed, any pending accounting information for each port meter created due to the definition of the meter in the policy is discarded.
Default	meter 1 (for unicast traffic) meter 9 multipoint (for all other traffic, other than unicast traffic)
Parameters	<i>meter-id</i> — Specifies the meter-id that uniquely identifies the meter within the policy. This is a required parameter each time the meter command is executed. Values 1 — 12 multipoint — This keyword specifies that this <i>meter-id</i> is for multipoint forwarded traffic only. This <i>meter-id</i> can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. If you attempt to map forwarding class unicast traffic to a multipoint queue, an error is generated and no changes are made to the current unicast traffic queue mapping. The meter must be created as multipoint. The multipoint designator cannot be defined after the meter is created. If an attempt is made to modify the command to include the multipoint keyword, an error is generated and the command will not execute. The multipoint keyword can be entered in the command line on a pre-existing multipoint meter to edit <i>meter-id</i> parameters. Values multipoint or not present Default Not present (the meter is created as non-multipoint)

meter

Syntax	meter <i>meter-id</i> no meter
Context	config>qos>network>ingress>fc
Description	This command overrides the default unicast forwarding type meter mapping for fc <i>fc-name</i>. The specified meter-id must exist within the policy as a non-multipoint meter before the mapping can be made. Once the forwarding class mapping is executed, all unicast traffic on a port using this policy is forwarded using the meter-id. The no form of this command sets the unicast (point-to-point) meter-id back to the default meter for the forwarding class (meter 1).
Default	meter 1
Parameters	<i>meter-id</i> — Specifies the meter-id. The specified parameter must be an existing, non-multipoint meter defined in the config>qos>network>ingress context. Values For network policy of type ip-interface: 1 — 12 (except 9, the default multipoint meter) For network policy of type port: 1 — 8

multicast-meter

Syntax	multicast-meter <i>meter-id</i> no multicast-meter
Context	config>qos>network>ingress>fc
Description	This command overrides the default multicast forwarding type meter mapping for fc <i>fc-name</i>. The specified meter-id must exist within the policy as a multipoint meter before the mapping can be made. Once the forwarding class mapping is executed, all multicast traffic on a port using this policy is forwarded using the meter-id. This command can only be used with a network policy of type ip-interface. The no form of the command sets the multicast forwarding type meter-id back to the default meter for the forwarding class.
Default	9
Parameters	<i>meter-id</i> — Specifies the multicast meter. The specified parameter must be an existing, multipoint meter defined in the config>qos>network>ingress context. Values 2 — 12

dscp

Syntax	dscp <i>dscp-name</i> fc <i>fc-name</i> profile { in out } no dscp <i>dscp-name</i>
Context	config>qos>network <i>policy-id</i> >ingress
Description	This command creates a mapping between the DiffServ Code Point (DSCP) of the network ingress traffic and the forwarding class. Ingress traffic that matches the specified DSCP will be assigned to the corresponding forwarding class. Multiple commands can be entered to define the association of some or all sixty-four DiffServ code points to the forwarding class. For undefined code points, packets are assigned to the forwarding class specified under the default-action command. The no form of this command removes the DiffServ code point to forwarding class association. The default-action then applies to that code point value.
Default	none
Parameters	<i>dscp-name</i> — The name of the DiffServ code point to be associated with the forwarding class. DiffServ code point can only be specified by its name and only an existing DiffServ code point can be specified. The software provides names for the well known code points. The system-defined names available are as follows. The system-defined names must be referenced as all lower case exactly as shown in the first column in Table 26 and Table 27 below. Additional names to code point value associations can be added using the 'dscp-name <i>dscp-name</i> <i>dscp-value</i>' command. The actual mapping is being done on the <i>dscp-value</i>, not the <i>dscp-name</i> that references the <i>dscp-value</i>. If a second <i>dscp-name</i> that references the same <i>dscp-value</i> is mapped within the policy, an error will occur. The second name will not be accepted until the first name is removed.

Table 26: Default DSCP Names to DSCP Value Mapping Table

DSCP Name	DSCP Value Decimal	DSCP Value Hexadecimal	DSCP Value Binary
nc1	48	0x30	0b110000
nc2	56	0x38	0b111000
ef	46	0x2e	0b101110
af41	34	0x22	0b100010
af42	36	0x24	0b100100
af43	38	0x26	0b100110
af31	26	0x1a	0b011010
af32	28	0x1c	0b011100
af33	30	0x1d	0b011110
af21	18	0x12	0b010010
af22	20	0x14	0b010100
af23	22	0x16	0b010110
af11	10	0x0a	0b001010
af12	12	0x0c	0b001100
af13	14	0x0e	0b001110
default	0	0x00	0b000000

Table 27: Default Class Selector Code Points to DSCP Value Mapping Table

DSCP Name	DSCP Value Decimal	DSCP Value Hexadecimal	DSCP Value Binary
cs7	56	0x38	0b111000
cs6	48	0x30	0b110000
cs5	40	0x28	0b101000
cs4	32	0x20	0b100000

Table 27: Default Class Selector Code Points to DSCP Value Mapping Table (Continued)

DSCP Name	DSCP Value Decimal	DSCP Value Hexadecimal	DSCP Value Binary
cs3	24	0x18	0b011000
cs2	16	0x10	0b010000
cs1	08	0x8	0b001000

fc *fc-name* — Enter this required parameter to specify the *fc-name* with which the code point will be associated.

Default none, for every DSCP value defined, the forwarding class must be indicated.

Values be, l2, af, l1, h2, ef, h1, nc

profile {**in** | **out**} — Enter this required parameter to indicate whether the DiffServ code point value is the in-profile or out-of-profile value.

NOTE 1: DSCP values mapping to forwarding classes Expedited (ef), High-1 (h1) and Network-Control (nc) can only be set to in-profile.

NOTE 2: DSCP values mapping to forwarding class ‘be’ can only be set to out-of-profile.

Default None, for every DSCP value defined, the profile must be indicated. If a DSCP value is not mapped, the default-action forwarding class and profile state will be used for that value.

Values in, out

lsp-exp

Syntax **lsp-exp** *lsp-exp-value* **fc** *fc-name* **profile** {**in** | **out**}
no lsp-exp *lsp-exp-value*

Context config>qos>network *policy-id*>ingress

Description This command creates a mapping between the LSP EXP bits of the network ingress traffic and the forwarding class.

Ingress traffic that matches the specified LSP EXP bits will be assigned to the corresponding forwarding class. Multiple commands can be entered to define the association of some or all eight LSP EXP bit values to the forwarding class. For undefined values, packets are assigned to the forwarding class specified under the **default-action** command.

The **no** form of this command removes the association of the LSP EXP bit value to the forwarding class. The **default-action** then applies to that LSP EXP bit pattern.

Default none

Parameters	<i>lsp-exp-value</i> — Specify the LSP EXP values to be associated with the forwarding class.
	Default None, the <i>lsp-exp</i> command must define a value.
	Values 0 to 8 (Decimal representation of three EXP bit field)
	fc fc-name — Enter this required parameter to specify the fc-name that the EXP bit pattern will be associated with.
	Default None, the <i>lsp-exp</i> command must define a fc-name.
	Values be, l2, af, l1, h2, ef, h1, nc
	profile {in out} — Enter this required parameter to indicate whether the LSP EXP value is the in-profile or out-of-profile value. The profile CLI parameter in the network qos policy of type ip-interface is deprecated.
	When 'no ldp-use-local-fc-enable' is set, the system will throw an warning and updates the associated mpls-lsp-exp-profile-map policy with new profile value specified by the user.
	When ldp-use-local-fc-enable is set, the system will error out the use of the profile command with an error message.
	Default None, the <i>lsp-exp</i> command must define a profile state.
	Values in, out

adaptation-rule

Syntax	adaptation-rule [<i>cir adaptation-rule</i>] [<i>pir adaptation-rule</i>] no adaptation-rule
Context	config>qos>network>ingress>meter
Description	This command defines the method used by the system to derive the operational CIR and PIR settings when the meter is provisioned in hardware. For the CIR and PIR parameters, individually the system attempts to find the best operational rate depending on the defined constraint. The no form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific adaptation-rule is removed, the default constraints for rate and cir apply.
Default	adaptation-rule cir closest pir closest
Parameters	<i>adaptation-rule</i> — Specifies the adaptation rule to be used while computing the operational CIR or PIR value. pir — Defines the constraints enforced when adapting the PIR rate defined within the meter <i>meter-id</i> rate command. The pir parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the meter. When the rate command is not specified, the default applies. cir — Defines the constraints enforced when adapting the CIR rate defined within the meter meter-id rate command. The cir parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the meter. When the cir parameter is not specified, the default constraint applies.

max — The **max** (maximum) option is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR/CIR will be the next multiple of 8 kbps that is equal to or lesser than the specified rate.

min — The **min** (minimum) option is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR/CIR will be the next multiple of 8 kbps that is equal to or higher than the specified rate.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR/CIR will be the next multiple of 8 kbps (that is closest to the specified rate.

cbs

Syntax **cbs** *size-in-kbits*
 no cbs

Context config>qos>network>ingress>meter

Description This command provides a mechanism to override the default reserved tokens for the meter. The committed burst size parameter specifies the maximum burst size that can be transmitted by the source while still complying with the CIR. If the transmitted burst is lower than the CBS value then the packets are marked as in-profile by the meter to indicate that the traffic is complying meter configured parameters.

The no form of this command returns the CBS size to the default value.

Default default

Parameters *size-in-kbits* — Specifies the size parameter is an integer expression of the number of kilobits reserved for the meter. For example, if a value of 10KBits is desired, then enter the value 10.

Values 324(for 7210 SAS E) — 16384, default
 For SAS-X: 4 — 2146959

mbs

Syntax **mbs** *size-in-kbits*
 no mbs

Context config>qos>network>ingress>meter

Description This command provides the explicit definition of the maximum amount of tokens allowed for a specific meter. The value is given in kilobits and overrides the default value for the context.

In case of trTCM, the maximum burst size parameter specifies the maximum burst size that can be transmitted by the source at the PIR while complying with the PIR. If the transmitted burst is lower than the MBS value then the packets are marked as out-profile by the meter to indicate that the traffic is not complying with CIR, but complying with PIR.

In case of srTCM, the maximum burst size parameter specifies the maximum burst size that can be transmitted by the source while not complying with the CIR. The transmitted burst is lower than the MBS value then the packets are marked as out-profile by the meter to indicate that the traffic is not complying with CIR.

If the packet burst is higher than MBS then packets are marked as red and dropped by the meter.

The **no** form of this command returns the MBS size assigned to the meter to the default value.

Default default

Parameters *size-in-kbits* — This parameter is an integer expression of the maximum number of kilobits of burst allowed for the meter. For example, for a value of 100 Kbits, enter the value 100.

Values 4 — 16384, default
For SAS-MX: 4 — 2146959

mode

Syntax **mode {trtcm | srtcm}**
no mode

Context config>qos>network>ingress>meter

Description This command defines the mode of the meter. The mode can be configured as Two Rate Three Color Marker (trTCM) or Single Rate Three Color Marker (srTCM). The mode command can be executed at anytime.

The **no** form of the command sets the default mode to be trtcm.

Default trtcm

Parameters **trtcm** — Meters the packet stream and marks the packets either green, yellow, or red. A packet is marked red if it exceeds the PIR. Otherwise, it is marked either yellow or green depending on whether it exceeds or doesn't exceed the CIR. The trTCM is useful, for example, for ingress policing of a service, where a peak rate needs to be enforced separately from a committed rate.

srtcm — Meters a packet stream and marks its packets either green, yellow, or red. Marking is based on a CIR and two associated burst sizes, a CBS and an Maximum Burst Size (MBS). A packet is marked green if it doesn't exceed the CBS, yellow if it does exceed the CBS, but not the cir and red otherwise. The srTCM is useful, for example, for ingress policing of a service, where only the length, not the peak rate, of the burst determines service eligibility.

rate

Syntax	rate cir <i>cir-rate-in-kbps</i> [pir <i>pir-rate-in-kbps</i>] no rate
Context	config>qos>network>ingress>meter
Description	This command defines the administrative PIR and CIR parameters for the meter. The rate command can be executed at anytime, altering the PIR and CIR rates for all meters created through the association of the Network QoS policy with the meter-id. The no form of the command returns all meter instances created with this meter-id to the default PIR and CIR parameters (max, 0).
Default	rate 0 pir max — The max default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The max value is mutually exclusive to the pir-rate value.
Parameters	cir <i>cir-rate-in-kbps</i> — The cir parameter overrides the default administrative CIR used by the meter. When the rate command has not been executed or the cir parameter is not explicitly specified, the default CIR (0) is assumed. Values 0 — , max pir <i>pir-rate-in-kbps</i> — Defines the administrative PIR rate, in kilobits, for the meter. When this rate command is executed, the PIR setting is optional. When the rate command has not been executed, the default PIR of max is assumed. Fractional values are not allowed and must be given as a positive integer. The actual PIR rate is dependent on the meter's adaptation-rule parameters and the actual hardware where the meter is provisioned. Values 1 — , max

Network Egress QoS Policy Commands

egress

Syntax **egress**

Context config>qos>network *policy-id*

Description This command is used to enter the CLI node that creates or edits egress policy entries that specify the forwarding class Dot1p marking map to be instantiated when this policy is applied to the network port.

The forwarding class and profile state mapping to EXP bits mapping for all packets are defined in this context.

All out-of-profile service packets are marked with the corresponding out-of-profile EXP bit value at network egress. All the in-profile service ingress packets are marked with the corresponding in-profile EXP bit value based on the forwarding class they belong.

fc

Syntax [**no**] **fc** *fc-name*

Context config>qos>network>egress

Description This command specifies the forwarding class name. The forwarding class name represents an egress queue. The **fc** *fc-name* represents a CLI parent node that contains sub-commands or parameters describing the egress characteristics of the queue and the marking criteria of packets flowing through it. The **fc** command overrides the default parameters for that forwarding class to the values defined in the network default policy. Appropriate default parameters are picked up based on whether the network-policy-type is port or ip-interface.

The **no** form of this command removes the forwarding class LSP EXP/Dot1p/DSCP associated with this , as appropriate. The forwarding class reverts to the defined parameters in the default network policy. If the *fc-name* is removed from the network policy that forwarding class reverts to the factory defaults.

Default Undefined forwarding classes default to the configured parameters in the default network policy policy-id 1.

Network Egress QoS Policy Commands

Parameters	<i>fc-name</i> — The case-sensitive, system-defined forwarding class name for which policy entries will be created.
Default	none
Values	be, l2, af, l1, h2, ef, h1, nc

Network Egress QoS Policy Forwarding Class Commands

fc

Syntax	[no] fc <i>fc-name</i> [create]
Context	config>qos>network>ingress
Description	This command creates a class instance of the forwarding class. Once the <i>fc-name</i> is created, classification actions can be applied and it can be used in match classification criteria. The no form of the command removes all the explicit meter mappings for <i>fc-name</i> forwarding types. The meter mappings revert to the default meters for <i>fc-name</i>.
Default	Undefined forwarding classes default to the configured parameters in the default policy <i>policy-id</i> 1.
Parameters	<i>fc-name</i> — The case-sensitive, system-defined forwarding class name for which policy entries will be created. Values be, l2, af, l1, h2, ef, h1, nc create — The keyword used to create the forwarding class. The create keyword requirement can be enabled/disabled in the environment>create context.

remark

Syntax	[no] remark < <i>policy-id</i> >
Context	config>qos>network <i>policy-id</i> >egress
Description	Allows the user to configure the remark policy to use marking the DSCP and Dot1p (802.1p) values in the packet header for IP packets sent out of this port. User should execute 'remarking' CLI command. If user enables remarking, without specifying a remark policy the system uses a default policy of remark-type 'dscp-only' Note: For an IP- Interface based policy, the remark policy type can be lsp-exp-only or dot1p-lsp-exp-shared. Note: For port based policy the remark policy can be dot1p-dscp.
Default	dscp-only
Parameters	<i>policy-id</i> — The policy ID of the remark policy.

remarking

Syntax **remarking**

Context config>qos>network *policy-id*>egress

Description This command remarks both access egress traffic and network egress traffic. The remarking is based on the forwarding class to LSP EXP/Dot1p/DSCP bit mapping defined under the egress node of the network QoS policy .

On network egress, for MPLS packets, only LSP EXP and Dot1p values can be marked. The LSP EXP mapping is defined in the network policy of type **ip-interface** and the Dot1p mapping can be defined in the network policy of type **port**.

On network egress, for IP packets, DSCP and Dot1p values can be marked. The Dot1p and DSCP values can be configured in the network policy of type **port**.

Normally, packets that ingress on network ports have, in case of MPLS packets, LSP EXP bit set by an upstream router. The packets are placed in the appropriate forwarding class based on the LSP EXP to forwarding class mapping. The LSP EXP bits of such packets are not altered as the packets egress this router, unless remarking is enabled.

Remarking can be required if this SAS X is connected to a different DiffServ domain where the EXP forwarding class mapping is different.

Typically, no remarking is necessary when all devices are in the same DiffServ domain. The network QoS policy supports an egress flag that forces remarking of packets that were received on network IP interfaces. This provides the capability of remarking without regard to the ingress state of the IP interface on which a packet was received. The effect of the setting of the egress network remark trusted state on each type of ingress IP interface and trust state is shown in the following table.

Ingress IP Interface Type and Trust State	Egress Network IP Interface Trust Remark Disabled (Default)	Egress Network IP Inter- face Trust Remark Enabled
Network Non-Trusted	Egress Remarked	Egress Remarked
Network Trusted (Default)	Egress Not Remarked	Egress Remarked

The remark trusted state has no effect on packets received on an ingress IP interface.

The remark trusted state is not applicable for network policies of type **port**.

The **no** form of this command reverts to the default behavior.

Default **no remarking** — Remarking disabled in the Network QoS policy.

Show Commands

network

Syntax	network [<i>policy-id</i>] [detail]
Context	show>qos
Description	This command displays network policy information.
Parameters	<i>policy-id</i> — Displays information for the specific policy ID. Defaultall network policies Values1 — 65535 detail — Includes information about ingress and egress bit mappings and network policy interface associations. Network QoS Policy Output Fields — The following table describes network QoS Policy output fields.

Table 28: Show QoS Network Output Fields

Label	Description
Policy-Id	The ID that uniquely identifies the policy.
Remark	True – Remarking is enabled for all packets that egress this router where the network policy is applied. The remarking is based on the forwarding class to bit mapping defined under the egress node of the network QoS policy.
Description	A text string that helps identify the policy's context in the configuration file.
Forward Class/ FC Name	Specifies the forwarding class name.
Profile	Out – In –
Accounting	Packet-based – Specifies that the meters associated with this policy do not account for packet framing overheads (such as Ethernet the Inter Frame Gap (IFG) and the preamble), while accounting for the bandwidth to be used by this flow. Frame-based – Specifies that the meters associated with this policy account for the packet framing overheads (such as for Ethernet the IFG and preamble), while accounting the bandwidth to be used by the flow.
Profile policy	Displays the profile policy ID.
Local FC	Displays if ldp-local-fc-enable is enabled or disabled
Global Prof	Displays the global profile policy ID for LDP packets.
Bit Mapping:	
Out-of-Profile	Displays the value used for out-of-profile traffic.
In-Profile	Displays the value used for in-profile traffic.
Interface	Displays the interface name.
IP Addr	Displays the interface IP address.
Port-Id	Specifies the physical port identifier that associates the interface.

*A:SAS-X-C>config>qos>network# show qos network 2 detail

Network Egress QoS Policy Forwarding Class Commands

```

=====
QoS Network Policy
=====
-----
Network Policy (2)
-----
Policy-id      : 2
Egr Remark     : False                      Egr Rem Plcy : N/A
Forward Class  : be                        Profile      : Out
Scope         : Template                  Policy Type   : IpInterface
Accounting     : packet-based
Profile Policy : 1
Local FC       : Disabled                  Global Prof   : 1
Description    : Default network QoS policy.

-----
Dot1p Bit Map          Forwarding Class          Profile
-----
No Matching Entries

-----
Meter Mode      CIR Admin CIR Rule    PIR Admin  PIR Rule    CBS Admin MBS Admin
                CIR Oper                                CBS Oper    MBS Oper
-----
1      TrTcm1_CA  0          closest    max        closest    def        def
                0          max
9      TrTcm1_CA  0          closest    max        closest    def        def
                0          max
-----

FC              UCastM          MCastM
-----
No FC-Map Entries Found.

-----
Interface Association
-----
Interface      : system
IP Addr.       : 180.10.10.10/32          Port Id       : system
Interface      : in-band-management
IP Addr.       : 10.135.25.183/24         Port Id       : 1/1/24
Interface      : management
IP Addr.       : 10.135.25.183/24         Port Id       : A/1
=====
*A:SAS-X-C>config>qos>network#

For SAS-MX:
*A:qos1# show qos network 1001 detail
=====
QoS Network Policy
=====
-----
Network Policy (1001)
-----
Policy-id      : 1001                      Remark        : False
Forward Class  : be                        Profile       : In
Attach Mode    : mpls                      Config Mode   : mpls

```

Scope : Template Policy Type : IpInterface
 Accounting : packet-based
 Description : ip-interface-type

LSP EXP Bit Map	Forwarding Class	Profile
0	be	Out
1	l2	Out
2	af	In
3	l1	Out
4	h2	In
5	ef	Out
6	h1	Out
7	nc	In

Meter Mode	CIR Admin CIR Oper	CIR Rule	PIR Admin PIR Oper	PIR Rule	CBS Admin CBS Oper	MBS Admin MBS Oper
1	TrTcm_CA	4000	closest	8000	closest	def
		4000		8000		def
2	TrTcm_CA	4000	closest	7000	closest	16384
		4000		7000		16000
3	TrTcm_CA	4000	closest	7000	closest	def
		4000		7000		def
4	TrTcm_CA	4000	closest	7000	closest	def
		4000		7000		def
5	TrTcm_CA	4000	closest	7000	closest	def
		4000		7000		def
6	TrTcm_CA	4000	closest	7000	closest	def
		4000		7000		def
7	TrTcm_CA	4000	closest	7000	closest	def
		4000		7000		def
8	TrTcm_CA	7000	closest	7000	closest	def
		7000		7000		def
9	TrTcm_CA	4000	closest	7000	closest	def
		4000		7000		def
10	TrTcm_CA	4000	closest	7000	closest	def
		4000		7000		def
11	TrTcm_CA	4000	closest	7000	closest	def
		4000		7000		def
12	TrTcm_CA	4000	closest	7000	closest	def
		4000		7000		def

FC	UCastM	MCastM
l2	2	def
af	3	def
l1	4	def
h2	5	12
ef	6	11
h1	7	10
nc	8	9

Egress Forwarding Class Queuing

FC Value : 0 FC Name : be
 - LSP EXP Bit Mapping
 Out-of-Profile : 0 In-Profile : 0

Network Egress QoS Policy Forwarding Class Commands

```
FC Value      : 1
- LSP EXP Bit Mapping
Out-of-Profile : 1
...
=====
*A:qos1#

*A:SAS-X-C>config>qos>network# show qos network 1 detail

=====
QoS Network Policy
=====
-----
Network Policy (1)
-----
Policy-id      : 1
Egr Remark     : False
Forward Class  : be
Scope          : Template
Accounting     : packet-based
Description    : Default network-port QoS policy.
Egr Rem Plcy   : N/A
Profile        : Out
Policy Type    : port

-----
Dot1p Bit Map          Forwarding Class          Profile
-----
No Matching Entries

-----
Meter Mode      CIR Admin CIR Rule   PIR Admin   PIR Rule   CBS Admin MBS Admin
                CIR Oper                PIR Oper                CBS Oper  MBS Oper
-----
1      TrTcm1_CA    0          closest    max        closest    def       def
                0                      max                      def       def

-----
FC              UCastM          MCastM
-----
No FC-Map Entries Found.

-----
Port Attachments
-----
Port-id : 1/1/10
Port-id : 1/1/11
Port-id : 1/1/12
Port-id : 1/1/13
Port-id : 1/1/20
Port-id : 1/1/21
Port-id : 1/1/22
Port-id : 1/1/23
Port-id : 1/1/24
Port-id : 1/1/25
Port-id : 1/1/26
Port-id : lag-3
Port-id : lag-5
=====
```

```
*A:SAS-X-C>config>qos>network#
```

mpls-lsp-exp-profile-map

- Syntax** `mpls-lsp-exp-profile-map [policy-id] [detail]`
- Context** `show>qos`
- Description** This command displays profile policy information.
- Parameters** *policy-id* — Displays information for the specific policy ID.
- Values** 1 — 65535
- detail** — Displays detail policy information.

Table 29: Show QoS Network Output Fields

Label	Description
Profile Map-id	Displays the profile Map ID.
Description	A text string that helps identify the policy's context in the configuration file.
Exp	Displays the EXP. values
Profile	Specifies the marking of the packets as in-profile or out-of-profile.
Network Policy Id	Displays the Network policy ID with which the mpls-lsp-exp-profile is associated.

Output

```
*A:7210-SAS>show>qos# mpls-lsp-exp-profile-map 1
```

```
=====
QoS MPLS LSP EXP Profile Maps
=====
-----
Profile Map-id      : 1
Description         : Default MPLS LSP EXP Profile Map policy
-----
Exp      Profile
-----
0        Out
1        In
2        Out
```

Network Egress QoS Policy Forwarding Class Commands

```
3      In
4      In
5      In
6      In
7      In
=====
*A:SAS-M>show>qos# mpls-lsp-exp-profile-map 1 detail
=====
QoS MPLS LSP EXP Profile Maps
=====
-----
Profile Map-id      : 1
Description         : Default MPLS LSP EXP Profile Map policy
-----
Exp      Profile
-----
0      Out
1      In
2      Out
3      In
4      In
5      In
6      In
7      In
-----
Network Policy Associations
-----
Network Policy Id      : 2
-----
=====
*A:7210-SAS>show>qos#
```

Network Queue QoS Policies

In This Section

This section provides information to configure network queue QoS policies using the command line interface.

Topics in this section include:

- [Overview on page 150](#)
- [Basic Configurations on page 152](#)
- [Default Network Queue Policy Values on page 155](#)
- [Service Management Tasks on page 158](#)

Overview

Network Queue policies define the egress network queuing for the traffic egressing on the network ports. Network queue policies are used at the Ethernet port and define the bandwidth distribution for the various FC traffic egressing on the Ethernet port.

There is one default network queue policy. Each policy always has 8 queues . Each of these queues are shared by unicast and multicast traffic. The default policies can be copied but they cannot be deleted or modified. The default policy is identified as **network-queue default**. Default network queue policies are applied to all network ports . You must explicitly create and then associate other network queue QoS policies.

Each queue has cir-level and pir-weight. Cir-level decides distribution of CIR traffic among queues while PIR-weight decides distribution of PIR traffic.

Cir-Level 8 has the highest priority and is serviced first(strict) irrespective of whether any other lower cir-levels have cir or not. For queue 8, it is recommended that CIR be set equal to PIR. When there are multiple queues in cir-level 1, they will share equal bandwidth. If there are multiple queues configured with cir-level 1 and some of them have CIR rate configured, then in the CIR loop, the bandwidth will be allocated to the queues that have CIR configured until all of the CIRs are satisfied and then in the PIR loop, the remaining bandwidth will be shared among all the queues. For example, if all queues have cir-level 1, and queue 1 and 2 have cir=100mbps, the throughput will be the following queues (considering the max speed of the port is 1-Gig):

- 1=200
- 2=200
- 3=100
- 4=100
- 5=100
- 6=100
- 7=100
- 8=100

Pir-Weight configured by the user is not considered by the system for a queue configured with cir-level 8. Therefore, if there are two queues with cir-level 8, all the traffic will be equally shared among the two, even when one queue have pir-weight 100 and other pir-weight as 1. pir-weight distributes the available bandwidth in the PIR loop among all egress network queues in the ratio of their weights.

For CIR Loop the priority between the cir-level are as follows:

8 , 7, 6, 5, 4, 3, 2 , 1

The system assigns the pir-level to the queues based on the cir-level configured by the user. There are 5 levels used by the system in the PIR loop. The assignment of PIR levels to CIR level is as shown in the [Table 30](#). For PIR Loop, queues at pir-level 5 is scheduled first, followed by queues at pir-level 4, and so on with queues at pir-level 1 being scheduled last.

Table 30: pir-level Assignment to queue based on the cir-level

Queue cir-level	Queue pir-level assigned
8	5
7	4
6,5	3
4,3	2
2,1	1

Configuration Guidelines

Queues at cir-level 8 are considered as strict priority queues by the system and therefore it is recommended that CIR value be set to PIR value. The system ignores the configured PIR value. Pir-level is not user configurable.

Basic Configurations

A basic network queue QoS policy must conform to the following:

- Each network queue QoS policy must have a unique policy name.
 - Queue parameters can be modified, but cannot be deleted.
-

Create a Network Queue QoS Policy

Configuring and applying QoS policies other than the default policy is optional. A default network queue policy is applied to all network ports.

To create an network queue policy, define the following:

- Enter a network queue policy name. The system will not dynamically assign a name.
- Include a description. The description provides a brief overview of policy features.
- FCs are mapped to 8 queues available at the port according to [Table 22, Forwarding Class to Queue-ID Map, on page 63](#).

Use the following CLI syntax to create a network queue QoS policy:

```
CLI Syntax: config>qos
             network-queue policy-name
                 description description-string
                 queue queue-id
                     rate cir cir-percent [pir pir-percent]
                     adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
                     [no] port-parent cir-level <cir-level 1-8> pir-weight
                           <pir-weight 1-100>
                     [no] queue-mgmt <queue-mgmt policy-name>
```

```
*A:SAS-X>config>qos>network-queue# info detail
```

```
-----
description "Default network queue QoS policy."
queue 1
  port-parent cir-level 1 pir-weight 1
  rate cir 0 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 2
  port-parent cir-level 2 pir-weight 1
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
```

```
exit
queue 3
  port-parent cir-level 3 pir-weight 1
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 4
  port-parent cir-level 4 pir-weight 1
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 5
  port-parent cir-level 5 pir-weight 1
  rate cir 100 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 6
  port-parent cir-level 6 pir-weight 1
  rate cir 100 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 7
  port-parent cir-level 7 pir-weight 1
  rate cir 10 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 8
  port-parent cir-level 8 pir-weight 1
  rate cir 10 pir 10
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
```

Applying Network Queue Policies

Apply network queue policies to the following entities:

- Ethernet Ports
-

Ethernet Ports

Use the following CLI syntax to apply a network queue policy to an Ethernet port.

CLI Syntax: config>port#
 ethernet
#-----
echo "Port Configuration"
#-----
 port 1/1/1
 ethernet
 mode network
 network
 queue-policy "nq1-cbs"
 exit
 exit
 exit
 no shutdown
exit

Default Network Queue Policy Values

The default network queue policies are identified as policy-id **default**. The default policies cannot be modified or deleted. The following displays default policy parameters:

```
*A:SAS-X>config>qos>network-queue# show qos network-queue default detail
```

```
=====
QoS Network Queue Policy
=====
-----
Network Queue Policy (default)
-----
Policy          : default
Accounting      : packet-based
Description     : Default network queue QoS policy.
-----
Queue Rates and Rules
-----
-----
QueueId         CIR(%)      CIR Adpt Rule   PIR(%)          PIR Adpt Rule
-----
Queue1          0             closest        100             closest
Queue2          25            closest        100             closest
Queue3          25            closest        100             closest
Queue4          25            closest        100             closest
Queue5          100           closest        100             closest
Queue6          100           closest        100             closest
Queue7          10            closest        100             closest
Queue8          10            closest        10             closest
-----
Parent Details
-----
-----
QueueId         Port        CIR Level      PIR Weight
-----
Queue1          True        1              1
Queue2          True        2              1
Queue3          True        3              1
Queue4          True        4              1
Queue5          True        5              1
Queue6          True        6              1
Queue7          True        7              1
Queue8          True        8              1
-----
High Slope
-----
-----
QueueId         State      Start-Avg(%)   Max-Avg(%)      Max-Prob(%)
-----
Queue1          Down      70             90              75
Queue2          Down      70             90              75
Queue3          Down      70             90              75
Queue4          Down      70             90              75
Queue5          Down      70             90              75
Queue6          Down      70             90              75
Queue7          Down      70             90              75
```

Default Network Queue Policy Values

Queue8	Down	70	90	75

Low Slope				

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)

Queue1	Down	50	75	75
Queue2	Down	50	75	75
Queue3	Down	50	75	75
Queue4	Down	50	75	75
Queue5	Down	50	75	75
Queue6	Down	50	75	75
Queue7	Down	50	75	75
Queue8	Down	50	75	75

Burst Sizes and Time Average Factor				

QueueId	CBS	MBS	Time Average Factor	

Queue1	def	def	7	
Queue2	def	def	7	
Queue3	def	def	7	
Queue4	def	def	7	
Queue5	def	def	7	
Queue6	def	def	7	
Queue7	def	def	7	
Queue8	def	def		

*A:SAS-X>config>qos>network-queue# info detail

```

-----
description "Default network queue QoS policy."
queue 1
  port-parent cir-level 1 pir-weight 1
  rate cir 0 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 2
  port-parent cir-level 2 pir-weight 1
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 3
  port-parent cir-level 3 pir-weight 1
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 4
  port-parent cir-level 4 pir-weight 1
  rate cir 25 pir 100
  adaptation-rule cir closest pir closest
  queue-mgmt "default"
exit
queue 5
  port-parent cir-level 5 pir-weight 1

```

```
        rate cir 100 pir 100
        adaptation-rule cir closest pir closest
        queue-mgmt "default"
    exit
queue 6
    port-parent cir-level 6 pir-weight 1
    rate cir 100 pir 100
    adaptation-rule cir closest pir closest
    queue-mgmt "default"
exit
queue 7
    port-parent cir-level 7 pir-weight 1
    rate cir 10 pir 100
    adaptation-rule cir closest pir closest
    queue-mgmt "default"
exit
queue 8
    port-parent cir-level 8 pir-weight 1
    rate cir 10 pir 10
    adaptation-rule cir closest pir closest
    queue-mgmt "default"
exit
-----
*A:SAS-X>config>qos>network-queue#
```

Service Management Tasks

This section discusses the following service management tasks:

- [Deleting QoS Policies on page 158](#)
 - [Copying and Overwriting QoS Policies on page 159](#)
 - [Editing QoS Policies on page 161](#)
-

Deleting QoS Policies

A network queue policy is associated by default with all network ports. You can replace the default policy with a customer-configured policy, but you cannot entirely remove a QoS policy. When you remove a QoS policy, the policy association reverts to the default network-queue policy **default**.

A network-queue policy cannot be deleted until it is removed from all network ports where it is applied.

To delete a user-created network queue policy, enter the following commands:

CLI Syntax: `config>qos# no network-queue policy-name`

Example: `config>qos# no network-queue nq1`

Copying and Overwriting QoS Policies

You can copy an existing network queue policy, rename it with a new policy ID name, or overwrite an existing network queue policy. The `overwrite` option must be specified or an error occurs if the destination policy ID exists.

CLI Syntax: `config>qos# copy network-queue source-policy-id dest-policy-id [overwrite]`

Example: `config>qos# copy network-queue nq1-cbs nq2-cbs`

The following output displays the copied policies

```
*A:card-1>config>qos# info
#-----
echo "QoS Slope and Queue Policies Configuration"
#-----
.....
    network-queue "nq1-cbs" create
        queue 1
            rate cir 0 pir 32
            adaptation-rule cir max
        exit
        queue 2
        exit
        queue 3
        exit
        queue 4
        exit
        queue 5
        exit
        queue 6
            rate cir 0 pir 4
        exit
        queue 7
            rate cir 3 pir 93
        exit
        queue 8
            rate cir 0 pir 3
        exit
    exit
network-queue "nq2-cbs" create
    queue 1
        rate cir 0 pir 32
        adaptation-rule cir max
    exit
    queue 2
    exit
    queue 3
    exit
    queue 4
    exit
    queue 5
```

Service Management Tasks

```
exit
queue 6
    rate cir 0 pir 4
exit
queue 7
    rate cir 3 pir 93
exit
queue 8
    rate cir 0 pir 3
exit
exit
-----
*A:card-1>config>qos# info
```

Editing QoS Policies

You can change existing policies, except the default policies, and entries in the CLI. The changes are applied immediately to all ports where the policy is applied. To prevent configuration errors use the copy command to make a duplicate of the original policy to a work area, make the edits, and then overwrite the original policy.

Network Queue QoS Policy Command Reference

Command Hierarchies

- [Configuration Commands on page 141](#)
- [Operational Commands on page 142](#)
- [Show Commands on page 142](#)

Configuration Commands

```
config
— qos
    — network-queue policy-name
        — description description-string
        — no description
        — queue queue-id [create]
            — adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
            — no adaptation-rule
            — port-parent [cir-level level] [cir-weight weight]
            — no port-parent
            — queue-mgmt <name>
            — no queue-mgmt
            — rate [cir cir-percent] [pir pir-percent]
            — no rate
```

Operational Commands

```
config
  — qos
    — copy network-queue src-name dst-name [overwrite]
```

Show Commands

```
show
  — qos
    — network-queue [network-queue-policy-name] [detail]
```

Configuration Commands

Generic Commands

description

Syntax **description** *description-string*
 no description

Context config>qos>network-queue

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Default No description is associated with the configuration context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands

copy

Syntax **copy network-queue** *src-name dst-name* [**overwrite**]

Context config>qos

Description This command copies or overwrites existing network queue QoS policies to another network queue policy ID.

The **copy** command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

Parameters **network-queue** — Indicates that the source policy ID and the destination policy ID are network-queue policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy.

overwrite — specifies to replace the existing destination policy. Everything in the existing destination policy will be overwritten with the contents of the source policy. If **overwrite** is not specified, a message is generated saying that the destination policy ID exists.

```
SR>config>qos# copy network-queue nq1 nq2
MINOR: CLI Destination "nq2" exists - use {overwrite}.
SR>config>qos# copy network-queue nq1 nq2 overwrite
```

Network Queue QoS Policy Commands

network-queue

Syntax	[no] network-queue <i>policy-name</i>
Context	config>qos
Description	This command creates a context to configure a network queue policy. Network queue policies on the Ethernet port define network egress queuing.
Default	default
Parameters	<i>policy-name</i> — The name of the network queue policy. Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Network Queue QoS Policy Queue Commands

queue

Syntax	queue <i>queue-id</i>
Context	config>qos>network-queue
Description	This command enables the context to configure a QoS network-queue policy queue. The FCs are mapped to these queues as per Table 22, Forwarding Class to Queue-ID Map, on page 63. Only one FC can be mapped to one queue. Queue-id 8 is the highest priority and Queue-id 1 is the lowest priority. Queue carry both the unicast and multicast traffic and no segregation is done. The hardware port scheduler prioritizes the queue according to the priority for each queue. High priority traffic should be mapped to high priority FC. Mapping traffic to high priority FC does not necessarily guarantee high priority treatment since the scheduler policy can influence the relative priority among the queues.
Parameters	<i>queue-id</i> — The <i>queue-id</i> for the queue, expressed as an integer. The <i>queue-id</i> uniquely identifies the queue within the policy. This is a required parameter each time the queue command is executed.

adaptation-rule

Syntax	adaptation-rule [<i>cir</i> <i>adaptation-rule</i>] [<i>pir</i> <i>adaptation-rule</i>] no adaptation-rule
Context	config>qos>network-queue>queue
Description	This command defines the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint. The no form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific adaptation-rule is removed, the default constraints for pir and cir apply.
Default	adaptation-rule cir closest pir closest
Parameters	<i>adaptation-rule</i> — Specifies the adaptation rule to be used while computing the operational CIR or PIR value. Values pir — Defines the constraints enforced when adapting the PIR rate defined within the queue queue-id rate command. The pir parameter requires a qualifier that defines the

constraint used when deriving the operational PIR for the queue. When the **pir** command is not specified, the default applies.

cir — Defines the constraints enforced when adapting the CIR rate defined within the **queue queue-id rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

max — The **max** (maximum) option is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) option is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue will be the rate closest to the rate specified using the **rate** command.

port-parent

Syntax	port-parent [cir-level <i>cir-level</i>] [pir-level <i>pir-weight</i>]
Context	config>qos>network-queue>queue
Description	This command specifies whether this queue feeds off a port-level scheduler. For the network-queue policy context, only the port-parent command is supported. It allows the user to configure the queues relative cir-level and pir-weight with respect to other queues on the port. The 7210 SAS X implements a port level scheduler that schedules all the queues associated with the port. The port level scheduler allocates bandwidth at line-rate or at the rate specified by the egress rate value configured for the port. The no form of the command sets the cir-level and pir-weight to default values.
Default	port-parent cir-level 1 pir-weight 1
Parameters	cir-level <i>cir-level</i> — Specifies the priority of the queue with respect to other queues. The priority of the queue is used only in the CIR loop. Level "8" is the highest priority and level "1" is the lowest priority. In the PIR loop, the priority of the queues cannot be configured. The system assigns the priority to the queues based on the cir-level associated with the queue. Values 1 — 8 (8 is the highest priority) Default 1

pir-weight *pir-weight* — Specifies the relative weight of the queue with respect to the other queues. The weight parameter is used only in the PIR loop. If a queue's level parameter is set to '8', the weight parameter is ignored by the system.

Values 1 — 100

Default 1

queue-mgmt

Syntax **queue-mgmt** < *name* >

Context config>qos>network-queue>queue

Description This command specifies the WRED and buffer parameters associated with the queue.

All the queues in the system allocate buffers from the system pool.

Parameters *name* — Specifies the name of the queue-management policy.

rate

Syntax **rate** [**cir** *cir-percent*] [**pir** *pir-percent*]
no rate

Context config>qos>network-queue>queue

Description This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the percentage at which the system prioritizes the queue over other queues competing for the same bandwidth. For network egress, the CIR also defines the rate that the queue is considered in-profile by the system. The in-profile and out-profile of the queue influences the scheduler priority queue metric. The in-profile and out-profile of the queue based on CIR and PIR is never marked in the packets. The packets at egress are considered in-profile and out-profile based on the SAP ingress policy meter results.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues. The 8 queues which are available at egress port are always associated with the network queue QoS policy by the queue-id.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (100, 0). (Except for the queue in cir-level 8, in which "no rate" sets cir=pir=1.)

If queue has cir-level as 8 then CIR should be equal to PIR for the queue, that is for queue with cir-level 8 all the traffic is considered as CIR(guaranteed).

- Parameters**
- cir percent** — Defines the percentage of the guaranteed rate allowed for the queue. When the **rate** command is executed, . When the **rate** command has not been executed, the default is assumed. Fractional values are not allowed and must be given as a positive integer.
- The actual rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.
- Values** 0 — 100
- Default** 0
- pir percent** — Defines the percentage of the maximum rate allowed for the queue. When the **rate** command is executed, the PIR setting is optional. When the **rate** command has not been executed, or the PIR parameter is not explicitly specified, the default PIR of 100 is assumed. Fractional values are not allowed and must be given as a positive integer.
- Values** 1 — 100 percent
- Default** 100

Show Commands

network-queue

- Syntax** **network-queue** [*network-queue-policy-name*] [**detail**]
- Description** This command displays network queue policy information.
- Context** show>qos
- Parameters** *network-queue-policy-name* — The name of the network queue policy.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

detail — Includes each queue’s rates and adaptation-rule and & cbs details. It also shows FC to queue mapping details.

Table 31: Network Queue Labels and Descriptions

Label	Description
Policy	The policy name that uniquely identifies the policy.
Description	A text string that helps identify the policy’s context in the configuration file.
Associations	Displays the physical port identifier where the network queue policy is applied.
Queue	Displays the queue ID.
CIR(%)	Displays the committed information rate
CIR Adapt Rule	Displays the adaptation rule in use
PIR(%)	Displays the peak information rate
PIR Adapt Rule	Displays the adaptation rule in use
Port	Indicates if the parent scheduler is port scheduler or not
CIR Level	Displays the priority of the queue in the CIR loop
PIR Weight	Displays the weight of the queue used in the PIR loop
High Slope	Displays the WRED high-slope parameters
Low Slope	Displays the WRED low-slope parameters

Table 31: Network Queue Labels and Descriptions

Label	Description
Burst Sizes (CBS/MBS)	Displays the configured CBS and MBS value
Time Avg Factor	Displays the WRED Time Average Factor value in use
FC and UcastQ	Displays the forwarding class (FC) to Queue association

```

*A:SAS-X>config>qos>network-queue# show qos network-queue default

=====
QoS Network Queue Policy
=====
-----
Network Queue Policy (default)
-----
Policy          : default
Accounting      : packet-based
Description     : Default network queue QoS policy.

-----
Associations
-----
Port-id : 1/1/10
Port-id : 1/1/19
Port-id : 1/1/20
Port-id : 1/1/24
=====

*A:SAS-X>config>qos>network-queue# show qos network-queue default detail

=====
QoS Network Queue Policy
=====
-----
Network Queue Policy (default)
-----
Policy          : default
Accounting      : packet-based
Description     : Default network queue QoS policy.

-----
Queue Rates and Rules
-----
-----
QueueId      CIR(%)    CIR Adpt Rule    PIR(%)    PIR Adpt Rule
-----
Queue1       0          closest         100       closest
Queue2       25         closest         100       closest
Queue3       25         closest         100       closest
Queue4       25         closest         100       closest
Queue5       100        closest         100       closest
Queue6       100        closest         100       closest
Queue7       10         closest         100       closest
Queue8       10         closest         10        closest
-----

```

Network Queue QoS Policy Commands

Parent Details

QueueId	Port	CIR Level	PIR Weight
Queue1	True	1	1
Queue2	True	2	1
Queue3	True	3	1
Queue4	True	4	1
Queue5	True	5	1
Queue6	True	6	1
Queue7	True	7	1
Queue8	True	8	1

High Slope

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	70	90	75
Queue2	Down	70	90	75
Queue3	Down	70	90	75
Queue4	Down	70	90	75
Queue5	Down	70	90	75
Queue6	Down	70	90	75
Queue7	Down	70	90	75
Queue8	Down	70	90	75

Low Slope

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	50	75	75
Queue2	Down	50	75	75
Queue3	Down	50	75	75
Queue4	Down	50	75	75
Queue5	Down	50	75	75
Queue6	Down	50	75	75
Queue7	Down	50	75	75
Queue8	Down	50	75	75

Burst Sizes and Time Average Factor

QueueId	CBS	MBS	Time Average Factor
Queue1	def	def	7
Queue2	def	def	7
Queue3	def	def	7
Queue4	def	def	7
Queue5	def	def	7
Queue6	def	def	7
Queue7	def	def	7
Queue8	def	def	7

Service Ingress QoS Policies

In This Section

This section provides information to configure SAP ingress QoS policies using the command line interface.

Topics in this section include:

- [Overview on page 176](#)
- [Basic Configurations on page 186](#)
- [Service Management Tasks on page 222](#)
- [Service Ingress Policy Configuration Considerations on page 180](#)
- [Allocation of QoS Resources for a SAP Ingress Policy on page 183](#)

Overview

There is one default service ingress policy. The default policy has only two meters.

The default policies can be copied but cannot be deleted. The default policies are identified as policy ID 1.

The default policies are applied to the appropriate interface, by default. For example, the default SAP ingress policy is applied to access ingress SAPs. You must explicitly associate other QoS policies.

For information about the tasks and commands necessary to access the command line interface and to configure and maintain your devices, refer to the CLI Usage chapter in the 7210 SR OS Basic System Configuration Guide.

Default SAP Ingress Policy

```
A:ALA-7>config>qos>sap-ingress$ info detail
-----
description "Default SAP ingress QoS policy"
scope template
meter 1 create
    mode trtcm
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    mbs default
    cbs default
exit
meter 11 multipoint create
    mode trtcm
    adaptation-rule cir closest pir closest
    rate cir 0 pir max
    mbs default
    cbs default
exit
default-fc be
-----
A:ALA-7>config>qos>sap-ingress$
```

SAP Ingress Policy Defaults

Table 32: SAP Ingress Policy Defaults

Field	Default
description	“Default SAP ingress QoS policy.”
scope	template
meter	1
mode	trtcm
adaptation-rule	cir closest pir closest
rate	pir = max, cir= 0
cbs	Default
mbs	Default
meter	11 (Multipoint)
mode	trtcm
adaptation-rule	cir closest pir closest
rate	pir = max, cir= 0
cbs	Default
mbs	Default
default-fc	be

Service Ingress Meter Selection Rules

The following are rules for meter selection by different traffic types under various configurations for VPLS services:

- If a FC is created without explicit meters, use the default meter 1 for unicast traffic and default meter 11 for all other traffic types (such as broadcast, multicast and unknownunicast).
- If a FC is created with an explicit unicast meter, use that meter for unicast traffic and use default meter 11 for all other traffic types.
- If a FC is created with an explicit unicast meter and explicit broadcast meter, use these meters for unicast and broadcast traffic respectively and use meter 11 for all other traffic types.
- If a FC is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for all other kinds of traffic.
- If a FC is created with an explicit unicast meter, an explicit broadcast meter, and an explicit multicast meter, use these meters for unicast, broadcast and multicast traffic types respectively. Unknown unicast traffic type will use the explicitly defined multicast meter.

The following are rules for meter selection for Epipe and VPRN services:

- A multipoint meter cannot be used. A multipoint meter configured in a policy is not used when the policy is applied to a SAP in an Epipe service.
- All FCs associated with a meter always use the unicast meter.

Service Ingress Policy Configuration Considerations

The *num-qos-classifiers* parameter cannot be modified when the policy is in use (for example, when it is associated with a SAP). Other parameters in the SAP ingress policy can be changed.

When changing other parameters (for example, fc meter map or fc classification match criteria) for a policy which is in use, the system recomputes the resources required due to accommodate the change. If the resources required exceeds the configured value for *num-qos-classifiers*, then the change is not allowed.

If more resources are needed than what is configured in *num-qos-classifiers* for an existing policy, then the following options are available.

- Copy the existing policy to a new policy, modify the *num-qos-classifiers* parameter, modify the match criteria entries suitably, and finally modify the SAP configuration to associate it with the new policy.
- Ensure the existing policy is not in use by any SAP (if required change the SAP configuration to disable the use of the QoS policy with the **no qos** form of the command), change all the required parameters and finally modify the SAP configuration to use the policy again.

Note that both these options have side-effects, for example, it resets the statistics associated with the meters and can potentially cause existing traffic classification not to take effect. But, the system will ensure that default policy is in use during the intermittent time when a policy change is being made following the steps given above.

- In releases prior to release 3.0R1, the software always computes the number of resources (like classifiers and meters) required by a policy assuming it will be used in a VPLS service. This allows the policy to be applied to either an Epipe or VPLS service.

From release 3.0R1 onwards, on creation of SAP ingress policy, software does not compute the number of resources required by a policy and validate it against resources available in the system. The software validates the resources needed only when the SAP ingress policy is attached to a SAP. If enough resources are available the SAP creation succeeds, else the software fails the CLI command.

Based on the service (i.e. Either VLL, VPLS, etc.) the SAP is configured in, for the same SAP ingress policy the amount of resources required is different. The software validates that the amount of qos resources specified with the command *num-qos-classifiers* is sufficient for the match criteria, forwarding class and service specified and the resources are available in hardware. On failure of the validation, the software disallows the association of the SAP ingress policy with the SAP.

For SAPs configured in VPRN services, the computation of resources is similar to an SAP configured in an Epipe service. Please see the section on "Allocation of QoS Resources for a SAP Ingress Policy" for more information.

SAP ingress hardware resources are allocated in chunks of fixed size (for 7210 SAS-M and 7210 SAS-X the chunk size is 512). A particular chunk can be used for either MAC criteria or IP criteria (unless if only dot1p bits or only DSCP bits are in use). In other words, a single chunk cannot be shared by both IP criteria and MAC criteria. If only dot1p bits are used in all the classification entries of the SAP ingress policy or only IP DSCP bits are used in all the classification entries of the SAP ingress policy, then the policy can use hardware resources from chunks currently in use by either IP-criteria or MAC-criteria. User can use the option 'dot1p-only' or 'dscp-only', if they plan to use only dot1p bits or only DSCP bits for SAP ingress classification. This typically allows for efficient use of available hardware resources and better scaling.

The system allocates the chunks on a first-cum-first-serve basis. In other words, user can create SAPs, with all of them using mac-criteria or user can create SAP with all of them using ip-criteria. When a SAP needs qos resources (i.e. when user associates a SAP ingress policy with the SAP) the system allocates the resources as follows:

- For SAPs configured in VPLS/EPIPE services (henceforth these SAPs are also referred to as L2 SAPs), the resources are always allocated from the mac-criteria chunks. If the resources required for the SAP is not available in one of the allocated chunks for mac-criteria, then the system allocates a new chunk from the shared pool if one is available. The remaining resource of the new chunk is available for use by mac-criteria policies of other SAPs. If no chunk is available in the shared pool, then the configuration command fails.
- For VPRN/L3 SAPs, if the SAP ingress policy is using ip-criteria, system checks if any of the existing chunks which are allocated to ip-criteria can accommodate the resources requested by the new SAP. If yes, the system allocates the resources from this chunk. If not, the system checks if a free chunk is available for use and allocates it. If the system cannot accommodate the new SAP in the existing chunks or no free chunks are available, the CLI command returns an error.
- For VPRN/L3 SAPs, if the SAP ingress policy is using dot1p-only or dscp-only, system checks if any of the existing chunks, which are in use for ip-criteria policies, can accommodate the resources requested by the new SAP. If yes, system allocates the resources from this chunk. If not, system checks if a free chunk is available for use and allocates it. The newly allocated chunk is accounted under ip-criteria if the SAP ingress policy which triggered the allocation specified an ip-criteria for classification or if doesn't have any criteria specified. If the system cannot accommodate the new SAP in the existing chunks or no free chunks are available, the CLI command returns an error.

Note: When the system allocates a new chunk, if the trigger is due to association of the SAP ingress policy to VPLS/EPIPE/L2 and the SAP ingress policy is configured to use dot1p-only criteria, then the remaining resources in the new chunk can be used for other SAPs using either dot1p-only policies, dscp-only policies or mac-criteria policies. Similarly, when a new chunk is

allocated, , if the trigger is due to association of the SAP ingress policy to VPRN/L3 SAP and the SAP ingress policy is configured to use dscp-only criteria, then the remaining resources in the new chunk can be used for other SAPs using either by dscp-only policies, dot1p-only policies or ip-criteria policies.

Note: The resource chunks referred to in this section is different from the resources specified using the command 'num-qos-classifiers'. The software manages the resource chunks. A SAP specifies the amount of qos resources it needs, using the 'num-qos-resources' CLI command (in the SAP ingress policy) and the system allocates the resources required by a SAP from the chunks depending on whether the SAP ingress policy uses ip-criteria or mac-criteria.

Allocation of QoS Resources for a SAP Ingress Policy

The user is allowed to configure the number of classification entries the SAP requires (for example: TQ).

Number of meters allocated automatically by system = $TQ / 2$ (up to a maximum of 32 meters).

To calculate the number of SAPs allowed, assume all configured to use 'TQ' QoS resources per SAP.

Number of SAPs allowed = maximum classification entries / TQ.

Note: If the number of SAPs is greater than the system limit, then the system limit takes precedence.

The user is allowed to mix and match SAPs with different QoS resources (that is, using different values of TQ).

The following determines the number of QoS resources to be allocated for an SAP:

- Number of match-criteria entries used to identify the FC.
- Number of FCs to use and number of traffic-types to be policed per FC.

Only those FCs that are in use by the match-criteria classification entries are considered for the number of FCs. Therefore, these FCs are referred to as 'FC in use'.

Given the number of traffic types to use per 'FC in use', the following rules apply for a SAP in a VPLS service to arrive at number of classification entries per FC in use:

- If a FC is in use and is created without explicit meters, use default meter #1 for unicast traffic and default meter #11 for all other traffic types (that is, broadcast, multicast and unknown-unicast). This requires two classification entries in hardware.
- If a FC is in use and is created with an explicit unicast meter, use that meter for unicast traffic and use default meter #11 for all other traffic types. This requires two classification entries in hardware.
- If a FC is in use and is created with an explicit unicast meter and explicit broadcast meter, use these meters for unicast and broadcast traffic respectively and use meter #11 for all other traffic types. This requires three classification entries in hardware.
- If a FC is in use and is created with an explicit unicast meter and explicit multicast meter, use the unicast meter for unicast traffic and multicast meter for all other kinds of traffic. This requires two classification entries in hardware.
- If a FC is in use and is created with an explicit unicast meter, an explicit broadcast meter, and an explicit multicast meter, use these meters for unicast, broadcast and multicast

traffic types respectively. Unknown unicast traffic type will use the explicitly defined multicast meter. This requires three classification entries in hardware.

For calculating the number of classification entries per FC for a SAP in a VLL or vprn service, the following rules apply:

- Multipoint meters cannot be used. Multipoint meter configured in a policy is not used when the policy is applied to a SAP in an Epipe service.
- All FCs in use and associated with a meter always use the unicast meter. Therefore, all FCs in use utilize only one classification entry in the hardware.

Apply the rules to determine the number of classification entries per FC (only for the FCs in use) using the following equation:

$$C(i) = \sum FCi(\text{unicast}) + FCi(\text{multicast}) + FCi(\text{broadcast}) + FCi(\text{unknown_unicast})$$

$$i = nc, h1, ef, h2, l1, af, l2, be$$

where $FCi(\text{unicast})$, $FCi(\text{multicast})$, $FCi(\text{broadcast})$, and $FCi(\text{unknown-unicast})$ are set to a value of 1 if this FC uses classifier to identify traffic-type unicast, multicast, broadcast and unknown-unicast respectively.

$FCi(\text{unicast})$, $FCi(\text{multicast})$, $FCi(\text{broadcast})$, and $FCi(\text{unknown-unicast})$ are set to a value of 0 if this FC does not use a classifier to identify this traffic-type.

If the user does not configure meters explicitly for the FC, then the default unicast meter and the default multicast meter are used. Therefore, by default, two classification entries in hardware are required by a FC.

Taking into account the number of match criteria and the number of FCs used, use the equation given below to arrive at total number of classification entries per policy, for example:

$$TC = \sum E(i) * C(i)$$

$$i = nc, h1, ef, h2, l1, af, l2, be$$

where:

- $E(i)$ is the number of match-criteria entries that classify packets to FCi . For 7210 platforms, the maximum number of classification entries per policy can be 64 (including default).
- $C(i)$ is the number of classification entries that are required by FCi to identify different traffic types.

Determine the number of policers or meters to use (for example TP). A maximum of 32 meters per policy are available.

Only those meters associated with FCs are considered for number of meters. Note that only 'FCs in use' is considered.

Total QoS resources required (for example TQ) = $\max((TC), (2 * TP))$.

The number obtained is rounded off to next binary number (power of 2).

The user configures value TQ using CLI command **num-qos-classifiers**.

Basic Configurations

A basic service ingress QoS policy must conform to the following:

- Have a unique service ingress QoS policy ID.
 - Have a QoS policy scope of template or exclusive.
 - Have at least one default unicast forwarding class .
 - Have at least one multipoint forwarding class .
-

Create Service Ingress QoS Policies

Configuring and applying QoS policies is optional. If no QoS policy is explicitly applied to a SAP, a default QoS policy is applied.

- [Service Ingress QoS Policy on page 187](#)

Service Ingress QoS Policy

To create an service ingress policy, define the following:

- A policy ID value. The system will not dynamically assign a value.
- Include a description. The description provides a brief overview of policy features.
- Specify *num-qos-classifiers* parameter. The default value is 16. The allowed values are 16, 36 and 256. The number of meters allowed is equal to half the number of QoS classifiers specified if the value is less than 32 or the number of meters is allowed is considered as 32.
- Specify a default forwarding class for the policy. All packets received on an ingress SAP using this ingress QoS policy will be classified to the default forwarding class.
- Define forwarding class parameters.
 - Modify the unicast-meter default value to override the default unicast forwarding type meter mapping for **fc** *fc-name*.
 - Modify the **multicast-meter** default value to override the default multicast forwarding type meters mapping for **fc** *fc-name*.
 - Modify the **unknown-meter** default value to override the default unknown unicast forwarding type **meter** mapping for **fc** *fc-name*.
 - Modify the **broadcast-meter** default value to override the default broadcast forwarding type **meter** mapping for **fc** *fc-name*.
- Specify IP or MAC criteria. You can define IP and MAC-based SAP ingress policies to select the appropriate ingress and corresponding forwarding class for matched traffic.
- A SAP ingress policy is created with a template scope. The scope can be modified to exclusive for a special one-time use policy. Otherwise, the **template** scope enables the policy to be applied to multiple SAPs.

The following displays an service ingress policy configuration:

```
A:ALA-7>config>qos>sap-ingress# info
-----
...
      sap-ingress 100 create
      description "Used on VPN sap"
...
-----
A:ALA-7>config>qos>sap-ingress#
```

Service Ingress QoS Meter

To create service ingress meter parameters, define the following:

- A new meter ID value — The system will not dynamically assign a value.
- Meter parameters — Ingress meters support the definition of either srTCM (Single Rate Tri-Color Meter) or trTCM (Two Rate Tri-Color Meter), CIR/PIR, CBS/MBS parameters.

The following displays an ingress meter configuration:

```
A:ALA-7>config>qos# info
#-----
echo "QoS Policy Configuration"
#-----
...
sap-ingress 100 create
  description "Used on VPN sap"
  meter 1 create
  exit
  meter 11 multipoint create
  exit
  meter 2 create
  rate cir 11000
  exit
  meter 3 create
  cbs 32
  rate 11000
  exit
  meter 4 create
  rate 1
  exit
  meter 5 create
  cbs 64
  mbs 128
  rate cir 1500 pir 1500
  exit
  meter 6 create
  mode srtcm
  rate cir 2500 pir 2500
  exit
  meter 7 multipoint create
  cbs 256
  mbs 512
  rate cir 100 pir 36
  exit
  meter 8 multipoint create
  cbs 256
  mbs 512
  rate cir 11000
  exit
  meter 9 multipoint create
  rate cir 11000
  exit
  meter 10 multipoint create
  rate cir 1
```

```
exit
meter 12 multipoint create
    rate cir 1500 pir 1500
exit
meter 13 multipoint create
    rate cir 2500 pir 2500
exit
meter 14 multipoint create
    rate cir 36 pir 100
exit
    meter 15 multipoint create
    rate cir 36 pir 100
exit
meter 16 multipoint create
    cbs 128
    mbs 256
    rate cir 36 pir 100
exit
...
#-----
A:ALA-7>config>qos#
```

SAP Ingress Forwarding Class (FC)

The following displays a forwarding class and precedence configurations:

```
A:ALA-7>config>qos# info
#-----
...
    fc af create
        meter 1
        broadcast-meter 7
        unknown-meter 8
    exit
    fc be create
        meter 2
        unknown-meter 9
    exit
    fc ef create
        meter 3
        broadcast-meter 10
    exit
    fc h1 create
        meter 4
        multicast-meter 12
    exit
    fc h2 create
        meter 5
        broadcast-meter 13
        multicast-meter 14
        unknown-meter 15
    exit
    fc nc create
        meter 6
        broadcast-meter 16
        multicast-meter 10
        unknown-meter 11
    exit
...
#-----
```

Service Ingress IP Match Criteria

When specifying SAP ingress match criteria, only one match criteria type can be configured in the SAP ingress QoS policy.

The following displays an ingress IP criteria configuration:

```
7210-SAS>config>qos>sap-ingress# info
-----
      num-qos-classifiers 32
      meter 1 create
      exit
      meter 11 multipoint create
      exit
      fc "h2" create
      exit
      ip-criteria any
        entry 16 create
        description "test"
        match
        exit
        action fc "be"
      exit
    exit
-----
7210-SAS>config>qos>sap-ingress#

7210-SAS>config>qos>sap-ingress# info
-----
      num-qos-classifiers 4
      meter 1 create
      exit
      meter 11 multipoint create
      exit
      ip-criteria dscp-only
        entry 30 create
        match
        exit
        action fc "l2"
      exit
    exit
-----
7210-SAS>config>qos>sap-ingress#
```

Service Ingress MAC Match Criteria

Both IP criteria and MAC criteria cannot be configured in the same SAP ingress QoS policy.

The MAC criteria "any" is only available for use with VPLS, VLL and PBB SAPs. MAC criteria "any" cannot be used with VPRN SAPs, only dot1p criteria can be used with VPRN SAPs.

To configure service ingress policy MAC criteria, define the following:

- A new entry ID value. Entries must be explicitly created. The system will not dynamically assign entries or a value.
- The action to associate the forwarding class with a specific MAC criteria entry ID.
- A description. The description provides a brief overview of policy features.

The following displays an ingress MAC criteria configuration:

```
7210-SAS>config>qos>sap-ingress# info
-----
description "test"
num-qos-classifiers 16
meter 1 create
exit
meter 11 multipoint create
exit
mac-criteria dot1p-only
  entry 25 create
    match
    exit
    no action
  exit
exit
default-fc "h1"
-----
7210-SAS>config>qos>sap-ingress#
```

Service Ingress QoS Policies Resource Usage Examples

The resource calculation shown for VLL is also applicable for VPRN services.

Example 1

```
sap-ingress 10 create
  description "example-policy-1"
  num-qos-classifiers 8
  meter 1
    rate cir 0 pir max
  exit
  meter 11 multipoint create
    rate cir 0 pir max
  exit
meter 3 create
  rate cir 100 pir 100
  exit
  scope template
  default-fcbe
  fc be
    meter 3
  exit
  fc af
    meter 1
  exit
  fc l1
    meter 3
  exit
  fc h2
    meter 3
  exit
  mac-criteria dot1p-only
    entry 1
      match dot1p 7
      action fc af
    exit
    entry 2
      match dot1p 5
      action fc l1
    exit
    entry 3
      match dot1p 6
      action fc h2
    exit
  exit
exit
```

In the example above, assuming the policy is attached to a SAP in a VPLS service, compute the number of classification entries per FC as follows:

```
FCnc = 0 + 0 + 0 + 0 = 0
FCh1 = 0 + 0 + 0 + 0 = 0
FCef = 0 + 0 + 0 + 0 = 0
```

Basic Configurations

$$FCh2 = 1 + 0 + 1 + 0 = 2$$

Since this FC uses unicast meter, need an entry to identify this traffic type explicitly. Another entry is needed to classify broadcast, multicast and unknown-unicast traffic type to the same FC and use the default meter #11.

$$FC11 = 1 + 0 + 1 + 0 = 2$$

$$FCaf = 1 + 0 + 1 + 0 = 2$$

$$FC12 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 1 + 0 = 2$$

Using the equation, calculate the total classification entries used by this policy, as follows:

$$TC = (0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (1 * 2)h2 + (1 * 2)l1 + (1 * 2)af + (0 * 0)l2 + (1 * 2)be = 8 \text{ (since three explicit match criteria entries are used to map traffic to each of FC H2, FC L1, and FC AF along with a default classification entry for FC BE).}$$

The total number of meters used = 3 (since FCs use meter #1, meter #3 and meter #11).

Hence, in this example, **num-qos-classifiers 8** is used (maximum of (8, (2 * 3))).

If the same policy were to be used for a SAP in an Epipe service, then since all traffic is classified to a unicast traffic type and since only unicast meters are used, the following:

$$FCnc = 0 + 0 + 0 + 0 = 0$$

$$FCh1 = 0 + 0 + 0 + 0 = 0$$

$$FCef = 0 + 0 + 0 + 0 = 0$$

$$FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FC11 = 1 + 0 + 0 + 0 = 1$$

$$FCaf = 1 + 0 + 0 + 0 = 1$$

$$FC12 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Using the above equation, total classification entries used = 4 and meters used = 2.

As can be seen here, using the same policy for Epipe SAP can lead to inefficient use of resources. Hence, it is recommended to create a different policy with the required number of resources (i.e. with **num-qos-classifiers 4**)

Example 2

```

sap-ingress 10 create
  description "example-policy-1"
  num-qos-classifiers 16
  meter 1
    rate cir 0 pir max
  exit
  meter 11 multipoint create
    rate cir 0 pir max
  exit
  meter 3 create
    rate cir 100 pir 100
  exit
  meter 2 multipoint create
    rate cir 1 pir 20
  exit
  scope template
  default-fcbe
  fc be
    meter 3
    broadcast-meter 2
  exit
  fc af
    meter 3
    broadcast-meter 2
  exit
  fc l1
    meter 3
    broadcast-meter 2
  exit
  fc h2
    meter 3
    broadcast-meter 2
  exit
  mac-criteria dot1p-only
entry 1
  match dot1p 7
  action fc af
exit
entry 2
  match dot1p 5
  action fc l1
exit
entry 3
  match dot1p 6
  action fc h2
exit
exit

```

In the example above, assuming the policy is attached to a SAP in a VPLS service, classification entries used per FC as:

```

FCnc = 0 + 0 + 0 + 0 = 0
FCh1 = 0 + 0 + 0 + 0 = 0
FCef = 0 + 0 + 0 + 0 = 0

```

Basic Configurations

$$FCh2 = 1 + 1 + 1 + 0 = 3$$

Since this FC uses unicast meter and broadcast meter, two entries are needed to identify these traffic types explicitly. Another entry is needed to classify multicast and unknown-unicast traffic type to the same FC and use the default meter #11.

$$FC11 = 1 + 1 + 1 + 0 = 3$$

$$FCaf = 1 + 1 + 1 + 0 = 3$$

$$FC12 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 1 + 1 + 0 = 3$$

Using the above equation, to get the total classification entries used = 12 (since three explicit match criteria entries map to each of FC H2, L1, and AF along with a default classification rule for BE).

The number of meters used = 3 (since FCs use only meter #2, meter #3 and meter #11).

Hence, in this example **num-qos-classifiers 16** is used (i.e. maximum of (12, (2*3))).

If the same policy were to be used for a SAP in an Epipe service, then since all traffic is classified to a unicast traffic type and since only unicast meters are used, the following:

$$FCnc = 0 + 0 + 0 + 0 = 0$$

$$FCh1 = 0 + 0 + 0 + 0 = 0$$

$$FCef = 0 + 0 + 0 + 0 = 0$$

$$FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FC11 = 1 + 0 + 0 + 0 = 1$$

$$FCaf = 1 + 0 + 0 + 0 = 1$$

$$FC12 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Using the above equation, to get total classification entries used = 4 and Meters used = 1.

As can be seen here, using the same policy for Epipe SAP can lead to inefficient use of resources. Hence, it is recommended to create a different policy with the required number of resources (i.e. with **num-qos-classifiers 4**)

Example 3

```

sap-ingress 10 create
  description "example-policy-2"
  num-qos-classifiers 16
  meter 1 create
    rate cir 100 pir 100
  exit
  meter 11 multipoint create
    rate cir 1 pir 20
  exit
  meter 3 create
    rate cir 100 pir 100
  exit
  meter 2 multipoint create
    rate cir 1 pir 20
  exit
  meter 4 multipoint create
    rate cir 10 pir 100
  exit
  meter 5 create
    rate cir 10 pir 10
  exit
  scope template
  default-fcbe
  fc af
    meter 3
    broadcast-meter 2
    multicast-meter 4
  exit
  fc l1
    meter 3
    broadcast-meter 2
  exit
  fc h2
    meter 3
    broadcast-meter 2
  exit
  fc h1
    meter 5
    broadcast-meter 4
    multicast-meter 4
    unknown-meter 4
  exit
  mac-criteria dot1p-only
    entry 1
      match dot1p 7
      action fc af
    exit
    entry 2
      match dot1p 5
      action fc l1
    exit
    entry 3
      match dot1p 6
      action fc h2
    exit
    entry 4

```

Basic Configurations

```
                match dot1p 3
                action fc h1
            exit
        exit
    exit
```

In the example above, assuming the policy is attached to a SAP in a VPLS service, the classification entries used per FC are:

$$\begin{aligned} \text{FCnc} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCh1} &= 1 + 1 + 1 + 1 = 4 \end{aligned}$$

Since this FC uses unicast, broadcast, multicast and unknown-unicast meter, four entries are needed to identify these traffic types explicitly.

$$\begin{aligned} \text{FCef} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCh2} &= 1 + 1 + 1 + 0 = 3 \end{aligned}$$

Since this FC uses unicast meter and broadcast meter, two entries are needed to identify these traffic types explicitly. Another entry is needed to classify multicast and unknown-unicast traffic to the same FC and use the default meter #11.

$$\text{FC11} = 1 + 1 + 1 + 0 = 3$$

Since this FC uses only unicast meter, an entry is needed to identify this traffic type explicitly. Another entry is needed to classify broadcast, multicast and unknown-unicast traffic type to the same FC and use the default meter #11.

$$\text{FCaf} = 1 + 1 + 1 + 0 = 3$$

Since this FC uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$\begin{aligned} \text{FC12} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCbe} &= 1 + 0 + 1 + 0 = 2 \end{aligned}$$

Using the above equation, the total classification entries used = 15 and meters used = 6.

If the same policy were to be used for a SAP in an Epipe service, then since all traffic is classified to a unicast traffic type and since only unicast meters are used, the following results:

$$\begin{aligned} \text{FCnc} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCh1} &= 1 + 0 + 0 + 0 = 1 \\ \text{FCef} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCh2} &= 1 + 0 + 0 + 0 = 1 \\ \text{FC11} &= 1 + 0 + 0 + 0 = 1 \\ \text{FCaf} &= 1 + 0 + 0 + 0 = 1 \\ \text{FC12} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCbe} &= 1 + 0 + 0 + 0 = 1 \end{aligned}$$

Using the above equation, the total classification entries used = 5 and meters used = 3 (since all FCs used only meter #1, meter #3 and meter #5).

Example 4

```
sap-ingress 10 create
  description"example-policy-3"
  num-qos-classifiers32
  meter 1 create
    rate cir100 pir 100
  exit
  meter11 multipoint create
    rate cir 1 pir 20
  exit
  meter 3 create
    rate cir100 pir 100
  exit
  meter2 multipoint create
    rate cir 1 pir 20
  exit
  meter 4 multipoint create
    rate cir 10 pir 100
  exit
  meter 5 create
    rate cir 10 pir 10
  exit
  meter 6 create
    rate cir 11 pir 100
  exit
  meter 8 multipoint create
    rate cir 20 pir 100
  exit
  scope template
  default-fcbe
  fc  af
    meter3
    broadcast-meter 2
    multicast-meter 4
  exit
  fc  l1
    meter3
    broadcast-meter 2
  exit
  fc  h2
    meter3
    broadcast-meter 2
  exit
  fc  h1
    meter 5
    broadcast-meter 4
    multicast-meter 4
    unknown-meter 4
  exit
  fc  ef
    meter 6
    broadcast-meter 2
    multicast-meter 8
  exit
  fc  nc
    meter 6
    broadcast-meter 2
```

```

        multicast-meter 8
    exit
    mac-criteria dot1p-only
        entry1
            match dot1p4
            action fc af
        exit
        entry 2
            match dot1p 5
            action fc l1
        exit
        entry 3
            match dot1p6
            action fc h2
        exit
        entry 4
            match dot1p 3
            action fc h1
        exit
        entry 5
            match dot1p 2
            action fc ef
        exit
        entry 6
            match dot1p 7
            action fc nc
        exit
    exit
exit

```

In the example above, assuming the policy is attached to a SAP in a VPLS service, compute the classification entries per FC as:

$$FCnc = 1 + 1 + 1 + 0 = 3$$

Since this FC uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCh1 = 1 + 1 + 1 + 1 = 4$$

Since this FC uses unicast, broadcast, multicast and unknown-unicast meter, four entries are needed to identify these traffic types explicitly.

$$FCef = 1 + 1 + 1 + 0 = 3$$

Since this FC uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCh2 = 1 + 1 + 1 + 0 = 3$$

Basic Configurations

Since this FC uses unicast meter and broadcast meter, two entries are needed to identify these traffic types explicitly. Another entry is needed to classify multicast and unknown-unicast traffic type to the same FC and use the default meter #11.

$$FC11 = 1 + 1 + 1 + 0 = 3$$

$$FCaf = 1 + 1 + 1 + 0 = 3$$

Since this FC uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FC12 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 1 + 0 = 2$$

Using the above equation, the total classification entries used = 21 and meters used = 8.

If the same policy were to be used for a SAP in an Epipe service, then since all traffic is classified to a unicast traffic type and since only unicast meters are used, the following:

$$FCnc = 1 + 0 + 0 + 0 = 1$$

$$FCh1 = 1 + 0 + 0 + 0 = 1$$

$$FCef = 1 + 0 + 0 + 0 = 1$$

$$FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FC11 = 1 + 0 + 0 + 0 = 1$$

$$FCaf = 1 + 0 + 0 + 0 = 1$$

$$FC12 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Using the above equation, the total classification entries used = 7 and meters used = 4.

As can be seen here, using the same policy for Epipe SAP can lead to inefficient use of resources. Hence, it is recommended to create a different policy with the required number of resources (i.e. with **num-qos-classifiers 8**)

Example 5

```

sap-ingress 10 create
  description "example-policy-3"
  num-qos-classifiers 32
  meter 1 create
    rate cir 100 pir 100
  exit
  meter 11 multipoint create
    rate cir 1 pir 20
  exit
  meter 3 create
    rate cir 100 pir 100
  exit
  meter 2 multipoint create
    rate cir 1 pir 20
  exit
  meter 4 multipoint create
    rate cir 10 pir 100
  exit
  meter 5 create
    rate cir 10 pir 10
  exit
  meter 6 create
    rate cir 11 pir 100
  exit
  meter 8 multipoint create
    rate cir 20 pir 100
  exit
  scope template
  default-fcbe
  fc af
    meter 3
    broadcast-meter 2
    multicast-meter 4
  exit
  fc l1
    meter 3
    broadcast-meter 2
  exit
  fc h2
    meter 3
    broadcast-meter 2
  exit
  fc h1
    meter 5
    broadcast-meter 4
    multicast-meter 4
    unknown-meter 4
  exit
  fc ef
  exit
  fc nc
    meter 6
    broadcast-meter 2
    multicast-meter 8
  exit
  mac-criteria dot1p-only

```

Basic Configurations

```
entry1
  match dot1p4
  action fc af
exit
entry 2
  match dot1p 5
  action fc l1
exit
entry 3
  match dot1p6
  action fc h2
exit
entry 4
  match dot1p 3
  action fc h1
exit
entry 5
  match dot1p 2
  action fc ef
exit
entry 6
  match dot1p 7
  action fc nc
exit
exit
```

In the example above, assuming the policy is attached to a SAP in a VPLS service, get the classification entries used per FC:

$$FCnc = 1 + 1 + 1 + 0 = 3$$

Since this FC uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FCh1 = 1 + 1 + 1 + 1 = 4$$

Since this FC uses unicast, broadcast, multicast and unknown-unicast meter, four entries are needed to identify these traffic types explicitly.

$$FCef = 1 + 0 + 1 + 0 = 2$$

Since no meters are explicitly configured, this FC uses the appropriate default meters all the traffic types (i.e. unicast traffic uses unicast meter #1 and broadcast, multicast, and unknown-unicast traffic uses multipoint meter #11).

$$FCh2 = 1 + 1 + 1 + 0 = 3$$

Since this FC uses unicast meter and broadcast meter, two entries are needed to identify these traffic types explicitly. Another entry is needed to classify multicast and unknown-unicast traffic to the same FC and use the default meter #11.

$$FCl1 = 1 + 1 + 1 + 0 = 3$$

$$FCaf = 1 + 1 + 1 + 0 = 3$$

Since this FC uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

$$FC12 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 1 + 0 = 2$$

Using the above equation, the total classification entries used = 20 and meters used = 8.

If the same policy were to be used for a SAP in an Epipe service, then since all traffic is classified to a unicast traffic type and since only unicast meters are used, the following:

$$FCnc = 1 + 0 + 0 + 0 = 1$$

$$FCh1 = 1 + 0 + 0 + 0 = 1$$

$$FCef = 1 + 0 + 0 + 0 = 1$$

$$FCh2 = 1 + 0 + 0 + 0 = 1$$

$$FC11 = 1 + 0 + 0 + 0 = 1$$

$$FCaf = 1 + 0 + 0 + 0 = 1$$

$$FC12 = 0 + 0 + 0 + 0 = 0$$

$$FCbe = 1 + 0 + 0 + 0 = 1$$

Using the above equation, to get the total classification entries used = 7 and meters used = 4.

Example 6

```
sap-ingress 10 create
  description "example-policy-1"
  num-qos-classifiers 16

  meter 1
    rate cir 0 pir max
  exit
  meter 11 multipoint create
    rate cir 0 pir max
  exit
  meter 3 create
    rate cir 100 pir 100
  exit
  meter 4 multipoint create
    rate cir 10 pir 50
  exit

  scope template

  default-fc      be
  fc  be
    meter 3
  exit
  fc  af
    meter 1
  exit
  fc  l1
    meter 3
    multicast-meter 4
  exit
  fc  h2
    meter 3
  exit

  mac-criteria dot1p-only
    entry 1
      match dot1p 7
      action fc af
    exit
    entry 2
      match dot1p 5
      action fc l1
    exit
    entry 3
      match dot1p 6
      action fc h2
    exit
  exit
exit
```

In the example above, assuming the policy is attached to a SAP in a VPLS service, the following number of classification entries per FC:

```
FCnc = 0 + 0 + 0 + 0 = 0
FCh1 = 0 + 0 + 0 + 0 = 0
```

$$\begin{aligned} \text{FCef} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCh2} &= 1 + 0 + 1 + 0 = 2 \\ \text{FC11} &= 1 + 0 + 1 + 0 = 2 \end{aligned}$$

Since this FC uses unicast meter and multicast meter, an entry is needed to identify these traffic types explicitly. Broadcast and unknown-unicast traffic is also classified using the same entry as multicast and use the same meter.

$$\text{FCaf} = 1 + 0 + 1 + 0 = 2$$

Since this FC uses unicast meter, an entry is needed to identify these traffic types explicitly. Another entry is needed to classify broadcast, multicast and unknown-unicast traffic type to the same FC and use the default meter #11.

$$\begin{aligned} \text{FC12} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCbe} &= 1 + 0 + 1 + 0 = 2 \end{aligned}$$

Using the above equation, the total classification entries used = 8 and meters used = 4.

If the same policy were to be used for a SAP in an Epipe service, then since all traffic is classified to a unicast traffic type and since only unicast meters are used, the following:

$$\begin{aligned} \text{FCnc} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCh1} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCef} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCh2} &= 1 + 0 + 0 + 0 = 1 \\ \text{FC11} &= 1 + 0 + 0 + 0 = 1 \\ \text{FCaf} &= 1 + 0 + 0 + 0 = 1 \\ \text{FC12} &= 0 + 0 + 0 + 0 = 0 \\ \text{FCbe} &= 1 + 0 + 0 + 0 = 1 \end{aligned}$$

Using the above equation, the total classification entries used = 4 and meters used = 2.

Example 7

```
sap-ingress 10 create
  num-qos-classifiers 8
  meter 1 create
  exit
  meter 11 multipoint create
  exit
  meter 3 create
  exit
  meter 4 multipoint create
  exit
  fc be
    meter 1
    broadcast-meter 11
    mulitcast-meter 4
  exit
  fc af
    meter 3
  exit
  default-fc be
  match entry 1
    dot1p 7 fc af
  exit
exit
```

In the example above, assuming the policy is attached to a SAP in a VPLS service, the following number of classification entries per FC are:

```
FCnc = 0 + 0 + 0 + 0 = 0
FCh1 = 0 + 0 + 0 + 0 = 0
FCef = 0 + 0 + 0 + 0 = 0
FCh2 = 0 + 0 + 0 + 0 = 0
FCl1 = 0 + 0 + 0 + 0 = 0
FCaf = 1 + 0 + 1 + 0 = 2
```

Since this FC uses unicast meter, an entry is needed to identify these traffic types explicitly. Another entry is needed entry to classify broadcast, multicast and unknown-unicast traffic type to the same FC and use the default meter #11.

```
FCl2 = 0 + 0 + 0 + 0 = 0
FCbe = 1 + 1 + 1 + 0 = 3
```

Since this FC uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

Using the above equation, the total classification entries used = 5 and meters used = 4.

If the same policy were to be used for a SAP in an Epipe service, then since all traffic is classified to a unicast traffic type and since only unicast meters are used, the following:

```
FCnc = 0 + 0 + 0 + 0 = 0
FCh1 = 0 + 0 + 0 + 0 = 0
```

$FCef = 0 + 0 + 0 + 0 = 0$
 $FCh2 = 0 + 0 + 0 + 0 = 0$
 $FC11 = 0 + 0 + 0 + 0 = 0$
 $FCaf = 1 + 0 + 0 + 0 = 1$
 $FC12 = 0 + 0 + 0 + 0 = 0$
 $FCbe = 1 + 0 + 0 + 0 = 1$

Using the above equation, the total classification entries used = 2 and meters used = 2.

Example 8

```
sap-ingress 10 create
  num-qos-classifiers16
  meter 1 create
  exit
  meter 11 multipoint create
  exit
  meter 3 create
  exit
  meter 4 multipoint create
  exit
  fc be
    meter 1
    broadcast-meter 11
    mulitcast-meter 4
  exit
  fc af
    meter 3
  exit
  default-fc be
  mac-criteria dot1p-only
  entry 1 create
    match dot1p 7 7
    action fc af
  exit
    dot1p 7 fc af
  exit
  match entry 2
    dot1p 5 fc af
  exit
  match entry 3
    dot1p 3 fc af
  exit
exit
```

In the example above, assuming the policy is attached to a SAP in a VPLS service, the following number of classification entries per FC:

```
FCnc = 0 + 0 + 0 + 0 = 0
FCh1 = 0 + 0 + 0 + 0 = 0
FCef = 0 + 0 + 0 + 0 = 0
FCh2 = 0 + 0 + 0 + 0 = 0
FC11 = 0 + 0 + 0 + 0 = 0
FCaf = 1 + 0 + 1 + 0 = 2
```

Since this FC uses unicast meter, an entry is needed to identify these traffic types explicitly. Another entry is needed to classify broadcast, multicast and unknown-unicast traffic type to the same FC and use the default meter #11.

```
FC12 = 0 + 0 + 0 + 0 = 0
FCbe = 1 + 1 + 1 + 0 = 3
```

Since this FC uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

Using the equation, calculate the total classification entries used by this policy, as follows

$$TC = (0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (0 * 0)h2 + (0 * 0)l1 + (3 * 2)af + (0 * 0)l2 + (1 * 3)be = 9$$

The number of meters used in this policy = 4.

Hence, in this example **num-qos-classifiers 16** is used (i.e. maximum of (9, (2 * 4))).

If the same policy were to be used for a SAP in an Epipe service, then since all traffic is classified to a unicast traffic type and since only unicast meters are used, the following:

$$\begin{aligned} FCnc &= 0 + 0 + 0 + 0 = 0 \\ FCh1 &= 0 + 0 + 0 + 0 = 0 \\ FCef &= 0 + 0 + 0 + 0 = 0 \\ FCh2 &= 0 + 0 + 0 + 0 = 0 \\ FC11 &= 0 + 0 + 0 + 0 = 0 \\ FCaf &= 1 + 0 + 0 + 0 = 1 \\ FC12 &= 0 + 0 + 0 + 0 = 0 \\ FCbe &= 1 + 0 + 0 + 0 = 1 \end{aligned}$$

Using the equation, calculate the total classification entries used by this policy, as follows:

$$(0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (0 * 0)h2 + (0 * 0)l1 + (3 * 1)af + (0 * 0)l2 + (1 * 1)be = 4$$

The number of meters used in this policy = 2.

Example 9

```
sap-ingress 10 create
  num-qos-classifiers 128
  meter 1 create
  exit
  meter 11 multipoint create
  exit
  meter 3 create
  exit
  meter 4 multipoint create
  exit
  fc be
    meter 1
    broadcast-meter 11
    mulitcast-meter 4
  exit
  fc af
    meter 3
    broadcast-meter 11
    multicast-meter 4
  exit
  default-fc be
  ip-criteria
  dscp-only
  match entry 1
    dscp 1 fc af
  exit
  match entry 2
    dscp 2 fc af
  exit
  match entry 3
    dscp 3 fc af
  exit
  match entry 4
    dscp 4 fc af
  exit
  match entry 5
    dscp 5 fc af
  exit
  match entry 6
    dscp 6 fc af
  exit
  match entry 7
    dscp 7 fc af
  exit
  match entry 8
    dscp 8 fc af
  exit
  match entry 9
    dscp 9 fc af
  exit
  match entry 10
    dscp 10 fc af
  exit
  match entry 11
    dscp 11 fc af
  exit
```

```
match entry 12
    dscp 12 fc af
exit
match entry 13
    dscp 12 fc af
exit
match entry 14
    dscp 14 fc af
exit
match entry 15
    dscp 15 fc af
exit
match entry 16
    dscp 16 fc af
exit
match entry 17
    dscp 17 fc af
exit
match entry 18
    dscp 18 fc af
exit
match entry 19
    dscp 19 fc af
exit
match entry 20
    dscp 20 fc af
exit
match entry 21
    dscp 21 fc af
exit
match entry 22
    dscp 22 fc af
exit
match entry 23
    dscp 23 fc af
exit
match entry 24
    dscp 24 fc af
exit
match entry 25
    dscp 25 fc af
exit
match entry 26
    dscp 26 fc af
exit
match entry 27
    dscp 27 fc af
exit
match entry 28
    dscp 28 fc af
exit
match entry 29
    dscp 29 fc af
exit
match entry 30
    dscp 30 fc af
exit
match entry 31
    dscp 31 fc af
```

Basic Configurations

```
exit
match entry 32
    dscp 32 fc af
exit
match entry 33
    dscp 33 fc af
exit
match entry 34
    dscp 34 fc af
exit
match entry 35
    dscp 35 fc af
exit
match entry 36
    dscp 36 fc af
exit
match entry 37
    dscp 37 fc af
exit
match entry 38
    dscp 38 fc af
exit
match entry 39
    dscp 39 fc af
exit
match entry 40
    dscp 40 fc af
exit
match entry 41
    dscp 41 fc af
exit
match entry 42
    dscp 42 fc af
exit
match entry 43
    dscp 43 fc af
exit
match entry 44
    dscp 44 fc af
exit
match entry 45
    dscp 45 fc af
exit
match entry 46
    dscp 46 fc af
exit
match entry 47
    dscp 47 fc af
exit
match entry 48
    dscp 48 fc af
exit
match entry 49
    dscp 49 fc af
exit
match entry 50
    dscp 50 fc af
exit
exit
```

exit

In the example above, assuming the policy is attached to a SAP in a VPLS service, the following number of classification entries per FC:

```
FCnc = 0 + 0 + 0 + 0 = 0
FCh1 = 0 + 0 + 0 + 0 = 0
FCef = 0 + 0 + 0 + 0 = 0
FCh2 = 0 + 0 + 0 + 0 = 0
FCl1 = 0 + 0 + 0 + 0 = 0
FCaf = 1 + 0 + 1 + 0 = 3
```

Since this FC uses unicast meter, an entry is needed to identify these traffic types explicitly. Another entry is needed to classify broadcast, multicast and unknown-unicast traffic type to the same FC and use the default meter #11.

```
FCl2 = 0 + 0 + 0 + 0 = 0
FCbe = 1 + 1 + 1 + 0 = 3
```

Since this FC uses unicast, broadcast and multicast meter, three entries are needed to identify these traffic types explicitly. Unknown-unicast traffic type is classified using the same entry as multicast traffic type and uses the same meter.

Using the equation, calculate the total classification entries used by this policy, as follows:

$$TC = (0 * 0)_{nc} + (0 * 0)_{h1} + (0 * 0)_{ef} + (0 * 0)_{h2} + (0 * 0)_{l1} + (50 * 3)_{af} + (0 * 0)_{l2} + (1 * 3)_{be} = 153$$

The number of meters used in this policy = 4.

Hence, in this example num-qos-classifiers 256 is used (maximum of (153, (2 * 4)) = 153, rounded off to the next binary power of 2 will be 256).

If the same policy were to be used for a SAP in an Epipe service, then since all traffic is classified to a unicast traffic type and since only unicast meters are used, e the following:

```
FCnc = 0 + 0 + 0 + 0 = 0
FCh1 = 0 + 0 + 0 + 0 = 0
FCef = 0 + 0 + 0 + 0 = 0
FCh2 = 0 + 0 + 0 + 0 = 0
FCl1 = 0 + 0 + 0 + 0 = 0
FCaf = 1 + 0 + 0 + 0 = 1
FCl2 = 0 + 0 + 0 + 0 = 0
FCbe = 1 + 0 + 0 + 0 = 1
```

Using the equation, calculate the total classification entries used by this policy, as follows:

$$(0 * 0)_{nc} + (0 * 0)_{h1} + (0 * 0)_{ef} + (0 * 0)_{h2} + (0 * 0)_{l1} + (50 * 1)_{af} + (0 * 0)_{l2} + (1 * 1)_{be} = 51$$

The number of meters used in this policy = 2.

Basic Configurations

Hence for Epipe SAP it is recommended to define another sap-ingress policy with num-qos-classifiers 64 is used (i.e. maximum of $(51, (2 * 2)) = 51$, rounded off to the next binary power of 2 will be 64).

Example 10

```

sap-ingress 10 create
  description "example-policy-1"
  num-qos-classifiers 4
  meter 1
    rate cir 0 pir max
  exit
  meter 11 multipoint create
    rate cir 0 pir max
  exit
  scope template
  default-fcl2
  fc l2
    meter1
  exit
  fc af
    meter1
  exit
  mac-criteria any
    entry1
      match dot1p7
      action fc af
    exit
  exit
exit

```

In the example above, assuming the policy is attached to a SAP in a VPLS service, compute the number of classification entries per FC as follows:

```

FCnc = 0 + 0 + 0 + 0 = 0
FCh1 = 0 + 0 + 0 + 0 = 0
FCef = 0 + 0 + 0 + 0 = 0
FCh2 = 0 + 0 + 0 + 0 = 0
FCl1 = 0 + 0 + 0 + 0 = 0
FCaf = 1 + 0 + 1 + 0 = 2
FCl2 = 1 + 0 + 1 + 0 = 2
FCbe = 0 + 0 + 0 + 0 = 2

```

Using the equation, calculate the total classification entries used by this policy, as follows:

$$TC = (0 * 0)_{nc} + (0 * 0)_{h1} + (0 * 0)_{ef} + (0 * 0)_{h2} + (0 * 0)_{l1} + (1 * 2)_{af} + (1 * 2)_{l2} + (0 * 0)_{be} = 4$$

The number of meters used = 2 (since both FCs use meter #1 and meter #11).

Hence, in this example **num-qos-classifiers 4** is used (i.e. maximum of (4, (2 * 2))).

If the same policy were to be used for a SAP in an Epipe service, then since all traffic is classified to a unicast traffic type and since only unicast meters are used, the following:

```

FCnc = 0 + 0 + 0 + 0 = 0
FCh1 = 0 + 0 + 0 + 0 = 0
FCef = 0 + 0 + 0 + 0 = 0
FCh2 = 0 + 0 + 0 + 0 = 0

```

Basic Configurations

```
FC11 = 0 + 0 + 0 + 0 = 0
FCaf = 1 + 0 + 0 + 0 = 1
FC12 = 1 + 0 + 0 + 0 = 1
FCbe = 0 + 0 + 0 + 0 = 0
```

Using the above equation, calculate the total classification entries used = 2 and meters used = 1.

As can be seen here, for Epipe SAP with the same amount of resources allocated one can have more FCs if need be.

Example 11

```

sap-ingress 10 create
  description "example-policy-1"
  num-qos-classifiers 4
  meter 1
    rate cir 0 pir max
  exit
  meter 11 multipoint create
    rate cir 0 pir max
  exit
  scope template
  default-fcbe
exit

```

In the example above, assuming the policy is attached to a SAP in a VPLS service, compute the number of classification entries per FC as follows:

```

FCnc = 0 + 0 + 0 + 0 = 0
FCh1 = 0 + 0 + 0 + 0 = 0
FCef = 0 + 0 + 0 + 0 = 0
FCh2 = 0 + 0 + 0 + 0 = 0
FCl1 = 0 + 0 + 0 + 0 = 0
FCaf = 0 + 0 + 0 + 0 = 0
FCl2 = 0 + 0 + 0 + 0 = 0
FCbe = 1 + 0 + 1 + 0 = 2

```

Using the equation, calculate the total classification entries used by this policy, as follows:

$$TC = (0 * 0)nc + (0 * 0)h1 + (0 * 0)ef + (0 * 0)h2 + (0 * 0)l1 + (0 * 0)af + (1 * 2)l2 + (0 * 0)be = 2$$

The number of meters used = 2 (since default FC uses meter #1 and meter #11).

Hence, in this example **num-qos-classifiers 4** is used (i.e. maximum of (2, (2 * 2))).

If the same policy were to be used for a SAP in an Epipe service, then since all traffic is classified to a unicast traffic type and since only unicast meters are used, the following:

```

FCnc = 0 + 0 + 0 + 0 = 0
FCh1 = 0 + 0 + 0 + 0 = 0
FCef = 0 + 0 + 0 + 0 = 0
FCh2 = 0 + 0 + 0 + 0 = 0
FCl1 = 0 + 0 + 0 + 0 = 0
FCaf = 0 + 0 + 0 + 0 = 0
FCl2 = 0 + 0 + 0 + 0 = 0
FCbe = 1 + 0 + 0 + 0 = 1

```

Using the above equation, total classification entries used = 1 and meters used = 1.

As can be seen here, for Epipe SAP with the same amount of resources allocated one can have more FCs if need be.

Applying Service Ingress Policies

Apply SAP ingress policies to the following service SAPs:

- [Epipe](#)
 - [VPLS](#)
 - [VPRN](#)
-

Epipe

The following output displays an Epipe service configuration with SAP ingress policy 100 applied to the SAP.

```
A:ALA-7>config>service# info
-----
      epipe 6 customer 6 vpn 6 create
        sap 1/1/18:3000 create
          ingress
            qos 100
          exit
        egress
          qos 200
        exit
      exit
    no shutdown
  exit
-----
A:ALA-7>config>service#
```

VPLS

The following output displays a VPLS service configuration with SAP ingress policy 100.

```
A:ALA-7>config>service# info
-----
      vpls 700 customer 7 vpn 700 create
        description "test"
        stp
          shutdown
        exit
      sap 1/1/9:10 create
        ingress
          qos 100
        exit
      exit
    exit
  exit
-----
```

```
A:ALA-7>config>service#
```

VPRN

The following output displays a VPRN service configuration.

```
A:ALA-7>config>service# info
-----
...
    vprn 1 customer 1 create
        ecmp 8
        autonomous-system 10000
        route-distinguisher 10001:1
        auto-bind ldp
        vrf-target target:10001:1
        interface "to-ce1" create
            address 11.1.0.1/24
            sap 1/1/10:1 create
                ingress
                    qos 100
                exit
                egress
                    qos 105
                exit
            exit
        exit
    exit
    no shutdown
exit
...
-----
A:ALA-7>config>service#
```

Service Management Tasks

This section discusses the following service management tasks:

- [Deleting QoS Policies on page 222](#)
 - [Copying and Overwriting QoS Policies on page 223](#)
 - [Remove a Policy from the QoS Configuration on page 224](#)
 - [Editing QoS Policies on page 224](#)
-

Deleting QoS Policies

Every service SAP is associated, by default, with the appropriate ingress policy (policy-id 1). You can replace the default policy with a customer-configured policy, but you cannot entirely remove the policy from the SAP configuration. When you remove a non-default service ingress policy, the association reverts to the default policy-id 1.

A QoS policy cannot be deleted until it is removed from all SAPs where they are applied.

```
A:ALA-7>config>qos# no sap-ingress 100
MINOR: CLI SAP ingress policy "100" cannot be removed because it is in use.
A:ALA-7>config>qos#
```

Remove a QoS Policy from Service SAP(s)

The following Epipe service output examples show that the SAP service ingress reverted to policy-id “1” when the non-default policies were removed from the configuration.

```
A:ALA-104>config>service>epipe# info detail
-----
      description "Distributed Epipe service to west coast"
      no tod-suite
      dot1ag
      exit
      ingress
        qos 1
        no filter
      exit
      egress
        no filter
      exit
      no collect-stats
      no accounting-policy
      no shutdown
-----
A:ALA-7>config>service>epipe#
```

Copying and Overwriting QoS Policies

You can copy an existing service ingress policy, rename it with a new policy ID value, or overwrite an existing policy ID. The `overwrite` option must be specified or an error occurs if the destination policy ID exists.

CLI Syntax: `config>qos# copy {sap-ingress } source-policy-id dest-policy-id [overwrite]`

```
*A:ALU-7210>config>qos# info
#-----
echo "QoS Policy Configuration"
#-----
      sap-ingress 100 create
        description "Used on VPN sap"
        meter 1 create
        exit
        meter 2 multipoint create
        exit
        meter 10 create
          rate cir 11000
        exit
        meter 11 multipoint create
        exit
      exit
      sap-ingress 101 create
        description "Used on VPN sap"
        meter 1 create
        exit
        meter 2 multipoint create
        exit
        meter 10 create
          rate cir 11000
        exit
        meter 11 multipoint create
        exit
      exit
      sap-ingress 200 create
        description "Used on VPN sap"
        meter 1 create
        exit
        meter 2 multipoint create
        exit
        meter 10 create
          rate cir 11000
        exit
        meter 11 multipoint create
        exit
      exit
#-----
*A:ALU-7210>config>qos#
```

Remove a Policy from the QoS Configuration

CLI Syntax: `config>qos# no sap-ingress policy-id`

Example: `config>qos# no sap-ingress 100`

Editing QoS Policies

You can change QoS existing policies and entries. The changes are applied immediately to all services where this policy is applied. To prevent configuration errors copy the policy to a work area, make the edits, and then write over the original policy.

Service SAP QoS Policy Command Reference

Command Hierarchies

- [Service Ingress QoS Policy Commands](#)
 - [FC Commands on page 235](#)
 - [MAC Criteria Commands on page 235](#)
- [Operational Commands](#)
- [Show Commands](#)

Service Ingress QoS Policy Commands

```
config
— qos
— [no] sap-ingress policy-id
— default-fc fc-name
— no default-fc
— description description-string
— no description
— [no] fc fc-name
— broadcast-meter meter-id
— no broadcast-meter
— meter meter-id
— no meter
— multicast-meter meter-id
— no multicast-meter
— unknown-meter meter-id
— no unknown-meter
— [no] ip-criteria [any | dscp-only]
— [no] entry entry-id
— action [fc fc-name]
— no action
— description description-string
— no description
— match [protocol protocol-id]
— no match
— dscp dscp-name
— no dscp
— dst-ip {ip-address/mask | ip-address netmask}
— no dst-ip
— dst-port {lt | gt | eq} dst-port-number
— dst-port range start end
— no dst-port
— fragment {true | false}
```

```

— no fragment
— src-ip {ip-address/mask | ip-address netmask}
— no src-ip
— src-port {lt | gt | eq} src-port-number
— src-port range start end
— no src-port
— renum [<old-entry-id> < new-entry-id>]
— [no] mac-criteria [any | dot1p-only]
— [no] entry entry-id
— action [fc fc-name]
— no action
— description description-string
— no description
— protocol-id: 0 — 255 protocol numbers accepted in DHB match
— no protocol-id: 0 — 255 protocol numbers accepted in DHB match
— dot1p dot1p-value [dot1p-mask]
— no dot1p
— dst-mac ieee-address [ieee-address-mask]
— no dst-mac
— etype etype-value
— no etype
— src-mac ieee-address [ieee-address-mask]
— no src-mac
— renum
— num-qos-classifiers [num-resources]
— meter meter-id [multipoint]
— no meter meter-id
— adaptation-rule [cir adaptation-rule] [pir adaptation-rule]
— no adaptation-rule
— cbs size-in-kbits
— no cbs
— mbs size-in-kbits
— no mbs
— mode {trtcn1 | trtcn2 | srtcn}
— no mode
— rate cir-rate-in-kbps [pir pir-rate-in-kbps]
— no rate
— scope {exclusive | template}
— no scope

```

Operational Commands

```
config
  — qos
    — copy sap-ingress src-pol dst-pol [overwrite]
```

Show Commands

```
show
  — qos
    — sap-ingress policy-id [detail]
```

Configuration Commands

Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>qos>sap-ingress config>qos>sap-ingress>ip-criteria>entry config>qos>sap-ingress>mac-criteria>entry
Description	This command creates a text description stored in the configuration file for a configuration context. The no form of this command removes any description string from the context.
Default	No description is associated with the configuration context.
Parameters	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands

copy

Syntax	copy sap-ingress <i>src-pol dst-pol</i> [overwrite]
Context	config>qos
Description	This command copies existing QoS policy entries for a QoS policy-id to another QoS policy-id. The copy command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the overwrite keyword.
Parameters	sap-ingress <i>src-pol dst-pol</i> — Indicates that the source policy ID and the destination policy ID are SAP ingress policy IDs. Specify the source policy ID that the copy command will attempt to copy from and specify the destination policy ID to which the command will copy a duplicate of the policy. Values 1 — 65535 overwrite — Specifies to replace the existing destination policy. Everything in the existing destination policy will be overwritten with the contents of the source policy. If overwrite is not specified, an error will occur if the destination policy ID exists.

renum

Syntax	renum <i>old-entry-number new-entry-number</i>
Context	config>qos>sap-ingress>ip-criteria config>qos>sap-ingress>mac-criteria
Description	This command renumbers existing QoS policy criteria entries to properly sequence policy entries. This can be required in some cases since the 7210 SAS exits when the first match is found and executes the actions in accordance with the accompanying action command. This requires that entries be sequenced correctly from most to least explicit.
Parameters	<i>old-entry-number</i> — Enter the entry number of an existing entry. Default none Values 1 — 4 <i>new-entry-number</i> — Enter the new entry-number to be assigned to the old entry. Default none Values 1 — 4

Service Ingress QoS Policy Commands

sap-ingress

Syntax [no] **sap-ingress** *policy-id*

Context config>qos

Description This command is used to create or edit the ingress policy. The ingress policy defines the Service Level Agreement (SLA) enforcement service packets receive as they ingress a SAP. SLA enforcement is accomplished through the definition of meters that have Forwarding Class (FC), Committed Information Rate (CIR), Peak Information Rate (PIR), Committed Burst Size (CBS), and Maximum Burst Size (MBS) characteristics. The simplest policy defines a single queue that all ingress traffic flows through. Complex policies have multiple meters combined with specific IP or MAC match criteria that indicate which queue a packet will flow through.

Policies in effect are templates that can be applied to multiple services as long as the **scope** of the policy is template. Meters defined in the policy are not instantiated until a policy is applied to a service SAP.

SAP ingress policies can be defined with either IP headers as the match criteria or MAC headers as the match criteria. The IP and MAC criteria are mutually exclusive and cannot be part of the same SAP ingress policy. Only one service ingress policy can be provisioned.

The SAP ingress policy with *policy-id* 1 is a system-defined policy applied to services when no other policy is explicitly specified. The system SAP ingress policy can be modified but not deleted. The **no sap-ingress** command restores the factory default settings when used on *policy-id* 1. The default SAP ingress policy defines one associated with the best effort (be) forwarding class, with CIR of zero and PIR of line rate.

Any changes made to the existing policy, using any of the sub-commands are applied immediately to all services where this policy is applied. For this reason, when many changes are required on a policy, it is recommended that the policy be copied to a work area policy ID. That work-in-progress policy can be modified until complete and then written over the original policy-id. Use the **config qos copy** command to maintain policies in this manner.

The **no sap-ingress policy-id** command deletes the SAP ingress policy. A policy cannot be deleted until it is removed from all services where it is applied. The system default SAP ingress policy is a special case; the **no** command restores the factory defaults to policy-id 1.

Parameters *policy-id* — The *policy-id* uniquely identifies the policy.

Values 1 — 65535

scope

Syntax	scope { exclusive template } no scope
Context	config>qos>sap-ingress <i>policy-id</i>
Description	This command configures the Service Ingress QoS policy scope as exclusive or template. The no form of this command sets the scope of the policy to the default of template .
Default	template
Parameters	exclusive — When the scope of a policy is defined as exclusive, the policy can only be applied to one SAP. If a policy with an exclusive scope is assigned to a second SAP an error message is generated. If the policy is removed from the exclusive SAP, it will become available for assignment to another exclusive SAP. template — When the scope of a policy is defined as template, the policy can be applied to multiple SAPs on the router. Default QoS policies are configured with template scopes. An error is generated when the template scope parameter to exclusive scope on default policies is modified.

default-fc

Syntax	default-fc <i>fc-name</i>
Context	config>qos>sap-ingress
Description	This command configures the default forwarding class for the policy. In the event that an ingress packet does not match a higher priority (more explicit) classification command, the default forwarding class will be associated with the packet. Unless overridden by an explicit forwarding class classification rule, all packets received on an ingress SAP using this ingress QoS policy will be classified to the default forwarding class. The default forwarding class is best effort (be). The default-fc settings are displayed in the show configuration and save output regardless of inclusion of the detail keyword.
Context	be
Parameters	<i>fc-name</i> — Specify the forwarding class name for the queue. The value given for <i>fc-name</i> must be one of the predefined forwarding classes in the system.

fc

Syntax	[no] fc <i>fc-name</i>
Context	config>qos>sap-ingress
Description	The fc command creates a class instance of the forwarding class <i>fc-name</i> . Once the <i>fc-name</i> is created, classification actions can be applied and can be used in match classification criteria.

The **no** form of the command removes all the explicit queue mappings for *fc-name* forwarding types. The queue mappings revert to the default meters for *fc-name*.

Parameters	<i>fc-name</i> — Specifies the forwarding class name for the queue. The value given for the <i>fc-name</i> must be one of the predefined forwarding classes for the system.		
Values	fc:	class	class: be, l2, af, l1, h2, ef, h1, nc
Default	None (Each class-name must be explicitly defined)		

ip-criteria

Syntax	[no] ip-criteria [any dscp-only]
Context	config>qos>sap-ingress
Description	IP criteria-based SAP ingress policies are used to select the appropriate ingress and corresponding forwarding class for matched traffic. User can specify either 'any' or 'dscp-only' as the sub-criteria. The sub-criteria determines what fields can be used to match traffic. The resource allocation for classification is affected by the sub-criteria in use. Please see the section on SAP ingress resource allocation for L2 and L3 criteria for more information. This command is used to enter the context to create or edit policy entries that specify IP criteria DiffServ code point. 7210 SAS OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason entries must be sequenced correctly from most to least explicit. The no form of this command deletes all the entries specified under this node. Once IP criteria entries are removed from a SAP ingress policy, the IP criteria is removed from all services where that policy is applied.
Default	dscp-only
Parameters	any — -Specifies that entries can use any of the fields available under ip-criteria (Example - IP source, IP destination, IP protocol fields can be used) dscp-only — Specifies that entries can use only the DSCP field.

mac-criteria

Syntax	[no] mac-criteria [any dot1p-only]
Context	config>qos>sap-ingress
Description	The mac-criteria based SAP ingress policies are used to select the appropriate ingress meters and corresponding forwarding class for matched traffic. User can specify either 'any' or 'dot1p-only' as the sub-criteria. The sub-criteria determines what fields can be used to match traffic. The resource allocation for classification is affected by the sub-criteria in use. Please

Service Ingress QoS Policy Commands

see the section on SAP ingress resource allocation for L2 and L3 criteria for more information.

This command is used to enter the node to create or edit policy entries that specify MAC criteria.

7210 SAS OS implementation will exit on the first match found and execute the actions in accordance with the accompanying action command. For this reason entries must be sequenced correctly from most to least explicit.

The **no** form of this command deletes all the entries specified under this node. Once mac-criteria entries are removed from a SAP ingress policy, the mac-criteria is removed from all services where that policy is applied.

Default any

Parameters **any** — Specifies that entries can use any of the fields available under mac-criteria (Example - MAC source, MAC destination, MAC Ethertype, etc. fields can be used)

dot1p-only — Specifies that entries can use only the Dot1p field.

num-qos-classifiers

Syntax **num-qos-classifiers** [*num-resources*]

Context config>qos>sap-ingress>num-qos-classifiers

Description This command configures the number of classifiers the SAP ingress Qos policy can use. A user cannot modify this parameter when it is in use (i.e. applied to a SAP).

The num-resources parameter also determines the maximum number of meters that are available to this policy. The maximum number of meters available for use by the forwarding classes (FC) defined under this policy is equal to half the value specified in the parameter num-resources. Any of these meters is available for use to police unicast or multipoint traffic. Any of these meters is available for use by more than one FC (or a single meter is available for use by all the FCs).

Default 4

Parameters *num-resources* — Specifies the number of resources planned for use by this policy

Values 4, 8, 16, 32, 64, 128, 256, 328

Service Ingress QoS Policy Forwarding Class Commands

broadcast-meter

Syntax	broadcast-meter <i>meter-id</i> no broadcast-meter
Context	config>qos>sap-ingress>fc
Description	This command overrides the default broadcast forwarding type meter mapping for fc <i>fc-name</i>. The specified <i>meter-id</i> must exist within the policy as a multipoint meter before the mapping can be made. Once the forwarding class mapping is executed, all broadcast traffic on a SAP using this policy will be forwarded using the <i>meter-id</i>. The broadcast forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior. The no form of the command sets the broadcast forwarding type <i>meter-id</i> back to the default of tracking the multicast forwarding type meter mapping.
Parameters	<i>meter-id</i> — Specifies an existing multipoint queue defined in the config>qos>sap-ingress context. Default 11

meter

Syntax	meter <i>meter-id</i> no meter
Context	config>qos>sap-ingress>fc
Description	This command overrides the default unicast forwarding type meter mapping for fc <i>fc-name</i>. The specified <i>meter-id</i> must exist within the policy as a non-multipoint meter before the mapping can be made. Once the forwarding class mapping is executed, all unicast traffic (this includes all traffic, even broadcast and multicast for services) on a SAP using this policy is forwarded using the <i>meter-id</i>. The no form of this command sets the unicast (point-to-point) <i>meter-id</i> back to the default meter for the forwarding class (meter 1).
Parameters	<i>meter-id</i> — Specifies an existing non-multipoint meter defined in the config>qos>sap-ingress context. Values 1 — 32

multicast-meter

Syntax	multicast-meter <i>meter-id</i> no multicast-meter				
Context	config>qos>sap-ingress>fc				
Context	This command overrides the default multicast forwarding type meter mapping for fc <i>fc-name</i> . The specified <i>meter -id</i> must exist within the policy as a multipoint meter before the mapping can be made. Once the forwarding class mapping is executed, all multicast traffic on a SAP using this policy is forwarded using the <i>meter-id</i> . The multicast forwarding type includes the unknown unicast forwarding type and the broadcast forwarding type unless each is explicitly defined to a different multipoint meter. When the unknown and broadcast forwarding types are left as default, they will track the defined meter for the multicast forwarding type. The no form of the command sets the multicast forwarding type <i>meter-id</i> back to the default meter for the forwarding class. If the broadcast and unknown forwarding types were not explicitly defined to a multipoint meter, they will also be set back to the default multipoint meter (11).				
Parameters	<i>meter-id</i> — Specifies an existing multipoint queue defined in the config>qos>sap-ingress context. <table> <tr> <td>Values</td><td>2— 32</td></tr> <tr> <td>Default</td><td>11</td></tr> </table>	Values	2— 32	Default	11
Values	2— 32				
Default	11				

unknown-meter

Syntax	unknown-meter <i>meter-id</i> no unknown-meter				
Context	config>qos>sap-ingress>fc				
Description	This command overrides the default unknown unicast forwarding type meter mapping for fc <i>fc-name</i> . The specified <i>meter-id</i> must exist within the policy as a multipoint meter before the mapping can be made. Once the forwarding class mapping is executed, all unknown traffic on a SAP using this policy is forwarded using the <i>meter-id</i> . The unknown forwarding type usually tracks the multicast forwarding type definition. This command overrides that default behavior. The no form of this command sets the unknown forwarding type <i>meter-id</i> back to the default of tracking the multicast forwarding type meter mapping.				
Parameters	<i>meter-id</i> — Specifies an existing multipoint meter defined in the config>qos>sap-ingress context. <table> <tr> <td>Values</td><td>2— 32</td></tr> <tr> <td>Default</td><td>11</td></tr> </table>	Values	2— 32	Default	11
Values	2— 32				
Default	11				

Service Ingress QoS Policy Entry Commands

action

Syntax	action [fc <i>fc-name</i>] no action
Context	config>qos>sap-ingress>ip-criteria>entry config>qos>sap-ingress>mac-criteria>entry
Description	This mandatory command associates the forwarding class with specific IP or MAC criteria entry ID. The action command supports setting the forwarding class parameter. Packets that meet all match criteria within the entry have their forwarding class overridden based on the parameters included in the action parameters. The action command must be executed for the match criteria to be added to the active list of entries. Each time action is executed on a specific entry ID, the previous entered values for fc <i>fc-name</i> is overridden with the newly defined parameters. The no form of the command removes the entry from the active entry list. Removing an entry on a policy immediately removes the entry from all SAPs using the policy. All previous parameters for the action is lost.
Default	Action specified by the default-fc .
Parameters	fc <i>fc-name</i> — The value given for fc <i>fc-name</i> must be one of the predefined forwarding classes in the system. Specifying the fc <i>fc-name</i> is required. When a packet matches the rule, the forwarding class is only overridden when the fc <i>fc-name</i> parameter is defined on the rule. If the packet matches and the forwarding class is not explicitly defined in the rule, the forwarding class is inherited based on previous rule matches.

entry

Syntax	[no] entry <i>entry-id</i> [create]
Context	config>qos>sap-ingress>ip-criteria config>qos>sap-ingress>mac-criteria
Description	This command is used to create or edit an IP or MAC criteria entry for the policy. Multiple entries can be created using unique <i>entry-id</i> numbers. The list of flow criteria is evaluated in a top down fashion with the lowest entry ID at the top and the highest entry ID at the bottom. If the defined match criteria for an entry within the list matches the information in the egress packet, the system stops matching the packet against the list and performs the matching entries reclassification actions. If none of the entries match the packet, the IP flow reclassification list has no effect on the packet. An entry is not populated in the list unless the action command is executed for the entry. An entry that is not populated in the list has no effect on egress packets. If the action command is executed without any explicit reclassification actions specified, the entry is populated in the list allowing packets matching the entry to

Service Ingress QoS Policy Entry Commands

exit the list, preventing them from matching entries lower in the list. Since this is the only flow reclassification entry that the packet matched and this entry explicitly states that no reclassification action is to be performed, the matching packet will not be reclassified.

The **no** form of this command removes the specified entry from the policy. Entries removed from the policy are immediately removed from all services where that policy is applied.

Default none

Parameters *entry-id* — The *entry-id*, expressed as an integer, uniquely identifies a match criterion and the corresponding action. It is recommended that multiple entries be given *entry-ids* in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

An entry cannot have any match criteria defined (in which case, everything matches) but must have at least the keyword **action fc** *fc-name* for it to be considered complete. Entries without the action keyword will be considered incomplete and hence will be rendered inactive.

Default none

Values 1 — 64

create — Required parameter when creating a flow entry when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the flow entry already exists.

match

Syntax [**no**] **match** [**protocol** *protocol-id*]

Context config>qos>sap-ingress>ip-criteria>entry

Description This command creates a context to configure match criteria for SAP QoS policy match criteria. When the match criteria have been satisfied the action associated with the match criteria is executed.

Only a single match criteria (either MAC or IP) is allowed at any point of time.

Parameters **protocol** *protocol-id* — Specifies an IP protocol to be used as a SAP QoS policy match criterion.

The protocol type such as TCP / UDP / OSPF is identified by its respective protocol number. Well-known protocol numbers include ICMP(1), TCP(6), UDP(17).

protocol-id:0 — 255 protocol numbers accepted in DHB

match

Syntax **match**
 no match

Context config>qos>sap-ingress>mac-criteria>entry

Description This command creates a context for entering/editing match MAC criteria for ingress SAP QoS policy match criteria. When the match criteria have been satisfied the action associated with the match criteria is executed.

If more than one match criteria (within one match statement) are configured then all criteria must be satisfied (AND function) before the action associated with the match will be executed.

A **match** context can consist of multiple match criteria, but multiple **match** statements cannot be entered per entry.

The **no** form of the command removes the match criteria for the *entry-id*.

IP QoS Policy Match Commands

dscp

Syntax	dscp no dscp
Context	config>qos>sap-ingress>ip-criteria>entry>match
Description	This command configures a DiffServ Code Point (DSCP) code point to be used for of packets from the specified FC. The no form of this command removes the DSCP match criterion.
Default	none
Parameters	<i>dscp-name</i> — Specifies a dscp name that has been previously mapped to a value using the dscp-name command. The DiffServ code point can only be specified by its name. Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, c p35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

dst-ip

Syntax	dst-ip { <i>ip-address/mask</i> <i>ip-address netmask</i> }
Context	config>qos>sap-ingress>ip-criteria>entry>match
Description	This command configures a destination address range to be used as a SAP QoS policy match criterion. To match on the destination address, specify the address and its associated mask, e.g., 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used. The no form of this command removes the destination IP address match criterion.
Default	none
Parameters	<i>ip-address</i> — The IP address of the destination IP or IPv6 interface. This address must be unique within the subnet and specified in dotted decimal notation. Values ip-address: a.b.c.d mask: 1 — 32 netmask a.b.c.d (dotted quad equivalent of mask length)

dst-port

Syntax	dst-port { lt gt eq } <i>dst-port-number</i> dst-port range <i>start end</i> no dst-port
Context	config>qos>sap-ingress config>qos>sap-ingress>ip-criteria>entry>match
Description	This command configures a destination TCP or UDP port number or port range for a SAP QoS policy match criterion. The no form of this command removes the destination port match criterion.
Default	none
Parameters	lt gt eq <i>dst-port-number</i> — The TCP or UDP port numbers to match specified as less than (lt), greater than (gt) or equal to (eq) to the destination port value specified as a decimal integer. Values 1 — 65535 (decimal hex or binary) range <i>start end</i> — The range of TCP or UDP port values to match specified as between the <i>start</i> and <i>end</i> destination port values inclusive. Values 1 — 65535 (decimal hex or binary)

fragment

Syntax	fragment { true false } no fragment
Context	config>qos>sap-ingress>ip-criteria>entry>match
Description	This command configures fragmented or non-fragmented IP packets as a SAP QoS policy match criterion. The no form of this command removes the match criterion.
Default	fragment false
Parameters	true — Configures a match on all fragmented IP packets. A match will occur for all packets that have either the MF (more fragment) bit set OR have the Fragment Offset field of the IP header set to a non-zero value. false — Configures a match on all non-fragmented IP packets. Non-fragmented IP packets are packets that have the MF bit set to zero and have the Fragment Offset field also set to zero.

src-ip

Syntax	src-ip <i>{ip-address/mask ip-address netmask}</i> no src-ip
Context	config>qos>sap-ingress>ip-criteria>entry>match config>qos>sap-egress>ip-criteria>entry>match config>qos>sap-ingress>ipv6-criteria>entry>match
Description	This command configures a source IP or IPv6 address range to be used as an SAP QoS policy match criterion. To match on the source IP or IPv6 address, specify the address and its associated mask, e.g. 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used. The no form of the command removes the source IP or IPv6 address match criterion.
Default	No source IP match criterion.
Parameters	<i>ip-address</i> <i>ipv6-address</i> — The IP or IPv6 address of the source IP interface. This address must be unique within the subnet and specified in dotted decimal notation. Values ip-address: a.b.c.d mask: 1 — 32 netmask a.b.c.d (dotted quad equivalent of mask length) <i>mask</i> — The subnet mask length, expressed as an integer or in dotted decimal notation. Values 0 — 32 <i>netmask</i> — Specify the subnet mask in dotted decimal notation. Values a.b.c.d (dotted quad equivalent of mask length)

src-port

Syntax	src-port <i>{lt gt eq} src-port-number</i> src-port range <i>start end</i> no src-port
Context	config>qos>sap-ingress>ip-criteria>entry>match
Description	This command configures a source TCP or UDP port number or port range for a SAP QoS policy match criterion. The no form of this command removes the source port match criterion.
Default	No src-port match criterion.
Parameters	lt gt eq <i>src-port-number</i> — The TCP or UDP port numbers to match specified as less than (lt), greater than (gt) or equal to (eq) to the source port value specified as a decimal integer. Values 1 — 65535 (decimal hex or binary)

range *start end* — The range of TCP or UDP port values to match specified as between the *start* and *end* source port values inclusive.

Values 1 — 65535 (decimal hex or binary)

Service Ingress MAC QoS Policy Match Commands

dot1p

Syntax **dot1p** *dot1p-value* [*dot1p-mask*]
no dot1p

Context config>qos>sap-ingress>mac-criteria>entry

Description The IEEE 802.1p value to be used as the match criterion.
 Use the **no** form of this command to remove the dot1p value as the match criterion.

Default None

Parameters *dot1p-value* — Enter the IEEE 802.1p value in decimal.

Values 0 — 7

dot1p-mask — This 3-bit mask can be configured using the following formats:

Format Style	Format Syntax	Example
Decimal	D	4
Hexadecimal	0xH	0x4
Binary	0bBBB	0b100

To select a range from 4 up to 7 specify *p-value* of 4 and a *mask* of 0b100 for value and mask.

Default 7 (decimal) (exact match)

Values 1 — 7 (decimal)

dst-mac

Syntax **dst-mac** *ieee-address* [*ieee-address-mask*]
no dst-mac

Context config>qos>sap-ingress>mac-criteria>entry

Description Configures a destination MAC address or range to be used as a Service Ingress QoS policy match criterion.
 The no form of this command removes the destination mac address as the match criterion.

Default none

Parameters *ieee-address* — The MAC address to be used as a match criterion.

Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

ieee-address-mask — A 48-bit mask to match a range of MAC address values.

This 48-bit mask can be configured using the following formats:

Format Style	Format Syntax	Example
Decimal	DDDDDDDDDDDDDD	281474959933440
Hexadecimal	0xHHHHHHHHHHHH	0xFFFFFFFF000000
Binary	0bBBBBBBB...B	0b11110000...B

All packets with a source MAC OUI value of 00-03-FA subject to a match condition should be specified as: 0003FA000000 0xFFFFFFFF000000

Default 0xFFFFFFFFFFFF (hex) (exact match)

Values 0x0000000000000000 — 0xFFFFFFFFFFFF (hex)

etype

Syntax **etype** *etype-value*
no etype

Context config>qos>sap-ingress>mac-criteria>entry

Description Configures an Ethernet type II value to be used as a service ingress QoS policy match criterion.

The Ethernet type field is a two-byte field used to identify the protocol carried by the Ethernet frame. For e.g. 0800 is used to identify the IP v4 packets.

The Ethernet type field is used by the Ethernet version-II frames. IEEE 802.3 Ethernet frames do not use the type field. For IEEE 802.3 frames use the dsap, ssap or snap-pid fields as match criteria.

The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and cannot be part of the same match criteria.

The no form of this command removes the previously entered etype field as the match criteria.

Default None

Parameters *etype-value* — The Ethernet type II frame Ethertype value to be used as a match criterion expressed in hexadecimal.

Values 0x0600 — 0xFFFF

src-mac

- Syntax

src-mac *ieee-address* [*ieee-address-mask*]
no src-mac
- Context

config>qos>sap-ingress>mac-criteria>entry
- Description

This command configures a source MAC address or range to be used as a service ingress QoS policy match criterion.

The **no** form of this command removes the source mac as the match criteria.
- Default

none
- Parameters

ieee-address — Enter the 48-bit IEEE mac address to be used as a match criterion.

Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

ieee-address-mask — This 48-bit mask can be configured using:

 This 48 bit mask can be configured using the following formats

Format Style	Format Syntax	Example
Decimal	DDDDDDDDDDDDDD	281474959933440
Hexadecimal	0xHHHHHHHHHHHH	0x0FFFFFF000000
Binary	0bBBBBBBB...B	0b11110000...B

To configure all packets with a source MAC OUI value of 00-03-FA are subject to a match condition, then the entry should be specified as: 003FA000000 0xFFFFFFFF000000

- Default

0xFFFFFFFFFFFF (hex) (exact match)
- Values

0x000000000000000 — 0xFFFFFFFFFFFF (hex)
- Values

Service Meter QoS Policy Commands

meter

Syntax **meter** *meter-id* [**multipoint**] [**create**]
 no meter *meter-id*

Context config>qos>sap-ingress

Description This command creates the context to configure an ingress service access point (SAP) QoS policy meter.

This command allows the creation of multipoint meters. Only multipoint meters can receive ingress packets that need flooding to multiple destinations. By separating the unicast for multipoint traffic at service ingress and handling the traffic on separate multipoint meters special handling of the multipoint traffic is possible. Each meter acts as an accounting and (optionally) policing device offering precise control over potentially expensive multicast, broadcast and unknown unicast traffic. Only the back-end support of multipoint traffic (between the forwarding class and the meter based on forwarding type) needs to be defined. The individual classification rules used to place traffic into forwarding classes are not affected. Meters must be defined as multipoint at the time of creation within the policy.

The multipoint meters are for multipoint-destined service traffic. Within non-multipoint services, such as Epipe services, all traffic is considered unicast due to the nature of the service type. Multicast and broadcast-destined traffic in an Epipe service will not be mapped to a multipoint service meter.

When an ingress SAP QoS policy with multipoint meters is applied to an Epipe SAP, the multipoint meters are not created.

Any billing or statistical queries about a multipoint meter on a non-multipoint service returns zero values. Any meter parameter information requested about a multipoint meter on a non-multipoint service returns the meter parameters in the policy. Multipoint meters would not be created for non-multipoint services.

The **no** form of this command removes the *meter-id* from the SAP ingress QoS policy and from any existing SAPs using the policy. Any forwarding class mapped to the meter, will revert to their default meters. When a meter is removed, any pending accounting information for each SAP meter created due to the definition of the meter in the policy is discarded.

Parameters *meter-id* — The *meter-id* for the meter, expressed as an integer. The *meter-id* uniquely identifies the meter within the policy. This is a required parameter each time the meter command is executed.

Values 1 — 32

adaptation-rule

Syntax **adaptation-rule** [**cir** *adaptation-rule*] [**pir** *adaptation-rule*]
 no adaptation-rule

Context config>qos>sap-ingress>meter

Description This command defines the method used by the system to derive the operational CIR and PIR settings when

the meter is provisioned in hardware. For the CIR and PIR parameters, individually the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific adaptation-rule is removed, the default constraints for **rate** and **cir** apply.

Default **adaptation-rule cir closest pir closest**

Parameters *adaptation-rule* — Specifies the adaptation rule to be used while computing the operational CIR or PIR value.

pir — Defines the constraints enforced when adapting the PIR rate defined within the meter meter-id rate command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the meter. When the rate command is not specified, the default applies.

cir — Defines the constraints enforced when adapting the CIR rate defined within the **meter rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the meter. When the **cir** parameter is not specified, the default constraint applies.

max — The **max** (maximum) option is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR/CIR will be the next multiple of 8 that is equal to or lesser than the specified rate.

min — The **min** (minimum) option is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR/CIR will be the next multiple of 8 kbps that is equal to or higher than the specified rate.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR/CIR will be the next multiple of 8 kbps that is closest to the specified rate.

cbs

Syntax **cbs size-in-kbits**
no cbs

Context config>qos>sap-ingress>meter

Description This command provides a mechanism to override the default CBS for the meter. The committed burst size parameter specifies the maximum burst size that can be transmitted by the source while still complying with the CIR. If the transmitted burst is lower than the CBS value then the packets are marked as in-profile by the meter to indicate that the traffic is complying meter configured parameters.

The **no** form of this command returns the CBS size to the default value.

Default default

Parameters *size-in-kbits* — Specifies the size parameter is an integer expression of the number of kilobits reserved for the meter. For example, if a value of 100 KBits is desired, then enter the value 100. The bucket size is rounded off to the next highest 4096 bytes boundary.

Values 4 — 16384, default
 For SAS-MX: 4-2146959

mbs

Syntax	mbs <i>size-in-kbits</i> no mbs
Context	config>qos>sap-ingress>meter
Description	This command provides the explicit definition of the maximum amount of tokens allowed for a specific meter. The value is given in Kilobits and overrides the default value for the context. In case of trtcm, the maximum burst size parameter specifies the maximum burst size that can be transmitted by the source at the PIR while complying with the PIR. If the transmitted burst is lower than the MBS value then the packets are marked as out-profile by the meter to indicate that the traffic is not complying with CIR, but complying with PIR. In case of srTCM, the MBS parameter specifies the maximum burst size that can be transmitted by the source while not complying with the CIR. The transmitted burst is lower than the MBS value then the packets are marked as out-profile by the meter to indicate that the traffic is not complying with CIR. If the packet burst is higher than MBS then packets are marked as red are dropped by the meter. The no form of this command returns the MBS size assigned to the meter to the value.
Default	default
Parameters	<i>size-in-kbits</i> — This parameter is an integer expression of the maximum number of Kilobits of buffering allowed for the meter. For example, for a value of 100 KBits, enter the value 100. Values 4 — 16384, default For SAS-MX: 4-2146959

mode

Syntax	mode {trtcm1 trtcm2 srtcm} no mode
Context	config>qos>sap-ingress>meter
Description	This command defines the mode of the meter. The mode can be configured as Two Rate Three Color Marker (trTCM1) or Single Rate Three Color Marker (srTCM). The mode command can be executed at anytime. Note: 1. The meter counters are reset to zero when the meter mode is changed. 2. For more information on the interpretation of rate parameters when the meter mode is configured as "trtcm2", refer to the command description of the policer rate command. The no form of the command sets the default mode trtcm1.
Default	trtcm1
Parameters	trtcm1 — Implements the policing algorithm defined in RFC2698. Meters the packet stream and marks its packets either green, yellow, or red. A packet is marked red if it exceeds the PIR. Otherwise, it is

marked either yellow or green depending on whether it exceeds or doesn't exceed the CIR. The trTCM1 is useful, for example, for ingress policing of a service, where a peak rate needs to be enforced separately from a committed rate. Two token buckets are used, the CBS bucket and the MBS bucket. Tokens are added to the buckets based on the CIR and PIR rates. The algorithm deducts tokens from both the CBS and the MBS buckets to determine a profile for the packet.

trtc2 — Implements the policing algorithm defined in RFC4115. Meters the packet stream and marks its packets either green, yellow, or red. A packet is marked red if it exceeds the PIR. Otherwise, it is marked either yellow or green depending on whether it exceeds or does not exceed the CIR. The trtc2 is useful, for example, for ingress policing of a service, where a peak rate needs to be enforced separately from a committed rate. Two token buckets are used, the CBS bucket and the EBS bucket. Tokens are added to the buckets based on the CIR and EIR rates. The algorithm deducts tokens from either the CBS bucket (that is, when the algorithm identifies the packet as in-profile or green packet) or the EBS bucket (that is, when the algorithm identifies the packet as out-of-profile or yellow packet).

Note: In this mode, the system configures the policer's EIR rate, based on the value of the PIR rate configured by the user.

srtcm — Meters an IP packet stream and marks its packets either green, yellow, or red. Marking is based on a CIR and two associated burst sizes, a CBS and an Maximum Burst Size (MBS). A packet is marked green if it doesn't exceed the CBS, yellow if it does exceed the CBS, but not the MBS, and red otherwise. The srTCM is useful, for example, for ingress policing of a service, where only the length, not the peak rate, of the burst determines service eligibility.

rate

Syntax	rate cir <i>cir-rate-in-kbps</i> [pir <i>pir-rate-in-kbps</i>] no rate
Context	config>qos>sap-ingress>meter
Description	This command defines the administrative PIR and CIR parameters for the meter. The rate command can be executed at anytime, altering the PIR and CIR rates for all meters created through the association of the SAP Ingress QoS policy with the meter-id. Note: When the meter mode is configured in trtc2 mode, the system interprets the PIR rate parameter as EIR for use by RFC 4115 algorithm. The no form of the command returns all meters created with the meter-id by association with the QoS policy to the default PIR and CIR parameters (max, 0).
Default	rate cir 0 pir max — The max default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The max value is mutually exclusive to the <i>pir-rate</i> value.
Parameters	cir <i>cir-rate-in-kbps</i> — The cir parameter overrides the default administrative CIR used by the meter. When the rate command has not been executed or the cir parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be given as a positive integer. The actual CIR rate is dependent on the meter's adaptation-rule parameters and the hardware. Values 0, max

pir *pir-rate-in-kbps* — Defines the administrative PIR rate, in kilobits, for the meter. When this command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of max is assumed. When the **rate** command is executed, a PIR setting is optional.

Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the meter's adaptation-rule parameters and the hardware.

Note: If the meter mode is configured as trtcm2, the system configures the policer's EIR rate, based on the value of the PIR rate configured by the user.

Values 0, max

Show Commands

sap-ingress

Syntax **sap-ingress** [*policy-id*] [**detail**]

Context show>qos

Description This command displays SAP ingress QoS policy information.

Parameters *policy-id* — Displays information about the specific policy ID.

Default all SAP ingress policies

Values 1 — 65535

detail — Displays detailed policy information including policy associations.

Sample Output

Show SAP Ingress Output — The following table describes SAP ingress show command output.

Label	Description
Policy-Id	The ID that uniquely identifies the policy.
Scope	Exclusive — Implies that this policy can only be applied to a single SAP. Template — Implies that this policy can be applied to multiple SAPs on the router.
Description	A text string that helps identify the policy's context in the configuration file.
Default FC	Specifies the default forwarding class for the policy.
Criteria-type	IP — Specifies that an IP criteria-based SAP ingress policy is used to select the appropriate ingress and corresponding forwarding class for matched traffic. MAC — Specifies that a MAC criteria-based SAP is used to select the appropriate ingress meters and corresponding forwarding class for matched traffic. Displays the meter ID.
Mode	Specifies the configured mode of the meter (trTcm1 or srTcm).

Label	Description (Continued)
CIR Admin	Specifies the administrative Committed Information Rate (CIR) parameters for the meters.
CIR Rule	<code>min</code> – The operational CIR for the meters will be equal to or greater than the administrative rate specified using the rate command. <code>max</code> – The operational CIR for the will be equal to or less than the administrative rate specified using the rate command. <code>closest</code> – The operational PIR for the meters will be the rate closest to the rate specified using the rate command without exceeding the operational PIR.
PIR Admin	Specifies the administrative Peak Information Rate (PIR) parameters for the meters.
PIR Rule	<code>min</code> – The operational PIR for the will be equal to or greater than the administrative rate specified using the rate command. <code>max</code> – The operational PIR for the meters will be equal to or less than the administrative rate specified using the rate command. <code>closest</code> – The operational PIR for the meters will be the rate closest to the rate specified using the rate command.
CBS	<code>def</code> – Specifies the default CBS value for the meters. <code>value</code> – Specifies the value to override the default reserved buffers for the meters.
MBS	<code>def</code> – Specifies the default MBS value. <code>value</code> – Specifies the value to override the default MBS for the .
UCast	Specifies the default unicast forwarding type meters mapping.
MCast	Specifies the overrides for the default multicast forwarding type mapping.
BCast	Specifies the default broadcast forwarding type meters mapping.
Unknown	Specifies the default unknown unicast forwarding type meters mapping.
Match Criteria Entry	Specifies an IP or MAC criteria entry for the policy.

Label	Description (Continued)
DSCP	Specifies a DiffServ Code Point (DSCP) name used for an ingress SAP QoS policy match.
FC	Specifies the entry's forwarding class.
Src MAC	Specifies a source MAC address or range to be used as a Service Ingress QoS policy match.
Dst MAC	Specifies a destination MAC address or range to be used as a Service Ingress QoS policy match.
Dot1p	Specifies a IEEE 802.1p value to be used as the match.
Ethernet-type	Specifies an Ethernet type II Ethertype value to be used as a Service Ingress QoS policy match.
FC	Specifies the entry's forwarding class.
Service Association	
Service-Id	The unique service ID number which identifies the service in the service domain.
Customer-Id	Specifies the customer ID which identifies the customer to the service.
SAP	Specifies the a Service Access Point (SAP) within the service where the SAP ingress policy is applied.
Classifiers required	Indicates the number of classifiers for a VPLS or Epipe service.
Meters required	Indicates the number of meters for a VPLS or Epipe service.

Sample Output

```
*A:Dut-G# show qos sap-ingress 100 detail
```

```
=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (100)
-----
Policy-id           : 100           Scope           : Template
Default FC         : be
Criteria-type      : MAC
Sub-Criteria-type   : dot1p-only
Accounting          : packet-based
Classifiers Allowed : 16           Meters Allowed    : 8
Classifiers Reqr'd (VPLS) : 16       Meters Reqr'd (VPLS) : 9 Exceeded
Classifiers Reqr'd (EPIPE) : 8       Meters Reqr'd (EPIPE) : 8
```

Description : (Not Specified)

Meter	Mode	CIR Admin	CIR Rule	PIR Admin	PIR Rule	CBS Admin	MBS Admin
1	TrTcm	0	closest	max	closest	def	def
2	TrTcm	0	closest	max	closest	def	def
3	TrTcm	0	closest	max	closest	def	def
4	TrTcm	0	closest	max	closest	def	def
5	TrTcm	0	closest	max	closest	def	def
6	TrTcm	0	closest	max	closest	def	def
7	TrTcm	0	closest	max	closest	def	def
8	TrTcm	0	closest	max	closest	def	def
9	TrTcm	0	closest	max	closest	def	def
10	TrTcm	0	closest	max	closest	def	def
11	TrTcm	0	closest	max	closest	def	def
12	TrTcm	0	closest	max	closest	def	def

FC	UCastM	MCastM	BCastM	UnknownM
l2	4	def	def	def
af	3	def	def	def
l1	5	def	def	def
h2	7	def	def	def
ef	2	def	def	def
h1	6	def	def	def
nc	8	def	def	def

Match Criteria

```

Entry                               : 1
Description                         : (Not Specified)
Src MAC                            :
Dst MAC                            :
Ethernet-type                       : Disabled
FC                                 : af

Entry                               : 2
Description                         : (Not Specified)
Src MAC                            :
Dst MAC                            :
Ethernet-type                       : Disabled
FC                                 : ef

Entry                               : 3
Description                         : (Not Specified)
Src MAC                            :
Dst MAC                            :
Ethernet-type                       : Disabled
FC                                 : l1

Entry                               : 4
Description                         : (Not Specified)
Src MAC                            :
Dst MAC                            :
Ethernet-type                       : Disabled
FC                                 : l2

Atm-Vci                            : Disabled
Dot1p                              : 1/7

Atm-Vci                            : Disabled
Dot1p                              : 2/7

Atm-Vci                            : Disabled
Dot1p                              : 3/7

Atm-Vci                            : Disabled
Dot1p                              : 4/7

```

Service Meter QoS Policy Commands

```

Entry                : 5
Description          : (Not Specified)
Src MAC              :                      Atm-Vci          : Disabled
Dst MAC              :                      Dot1p            : 5/7
Ethernet-type        : Disabled
FC                   : h1

```

```

Entry                : 6
Description          : (Not Specified)
Src MAC              :                      Atm-Vci          : Disabled
Dst MAC              :                      Dot1p            : 6/7
Ethernet-type        : Disabled
FC                   : h2

```

```

Entry                : 7
Description          : (Not Specified)
Src MAC              :                      Atm-Vci          : Disabled
Dst MAC              :                      Dot1p            : 7/7
Ethernet-type        : Disabled
FC                   : nc

```

----- Associations

```

Service-Id           : 100 (Epipe)      Customer-Id         : 1
- SAP : 1/1/1:100

```

```

=====
*A:Dut-G#

```

```

*A:qos1# show qos sap-ingress 102 detail

```

=====

QoS Sap Ingress

=====

Sap Ingress Policy (102)

```

-----
Policy-id            : 102                      Scope              : Template
Default FC           : be
Criteria-type        : MAC
Sub-Criteria-type    : dot1p-only
Accounting           : packet-based
Classifiers Allowed: 32                      Meters Allowed      : 16
Classifiers Used     : 32                      Meters Used         : 16

```

```

-----
Meter Mode  CIR Admin  CIR Rule  PIR Admin  PIR Rule  CBS  MBS
-----
1   TrTcm    100      closest    200        closest   32   128
2   TrTcm    100      closest    200        closest   32   128
3   TrTcm    100      closest    200        closest   32   128
4   TrTcm    100      closest    200        closest   32   128
5   TrTcm    100      closest    200        closest   32   128
6   TrTcm    100      closest    200        closest   32   128
7   TrTcm    100      closest    200        closest   32   128
8   TrTcm    100      closest    200        closest   32   128
9   TrTcm    100      closest    200        closest   32   128
10  TrTcm    100      closest    200        closest   32   128
11  TrTcm    100      closest    200        closest   32   128
12  TrTcm    100      closest    200        closest   32   128

```

13	TrTcm	100	closest	200	closest	32	128
14	TrTcm	100	closest	200	closest	32	128
15	TrTcm	100	closest	200	closest	32	128
16	TrTcm	100	closest	200	closest	32	128

FC		UCastM		MCastM		BCastM	
						UnknownM	

be		1		11		16	
l2		8		16		16	
af		7		15		16	
l1		6		14		16	
h2		5		13		16	
ef		4		12		16	
h1		3		10		16	
nc		2		9		16	

Match Criteria							

Entry	:	1					
Src MAC	:						
Dst MAC	:					Dot1p	: 7/7
Ethernet-type	:	Disabled					
FC	:	nc					
Entry	:	2					
Src MAC	:						
Dst MAC	:					Dot1p	: 6/7
Ethernet-type	:	Disabled					
FC	:	h1					
Entry	:	3					
Src MAC	:						
Dst MAC	:					Dot1p	: 5/7
Ethernet-type	:	Disabled					
FC	:	ef					
Entry	:	4					
Src MAC	:						
Dst MAC	:					Dot1p	: 4/7
Ethernet-type	:	Disabled					
FC	:	h2					
Entry	:	5					
Src MAC	:						
Dst MAC	:					Dot1p	: 3/7
Ethernet-type	:	Disabled					
FC	:	l1					
Entry	:	6					
Src MAC	:						
Dst MAC	:					Dot1p	: 2/7
Ethernet-type	:	Disabled					
FC	:	af					
Entry	:	7					
Src MAC	:						
Dst MAC	:					Dot1p	: 1/7
Ethernet-type	:	Disabled					
FC	:	l2					

Service Meter QoS Policy Commands

```

-----
Associations
-----
Service-Id      : 102 (VPLS)           Customer-Id      : 1
- SAP : 1/1/3:102
- SAP : 1/1/7:102
=====
*A:qos1#

```

For SAS-MX:

```
*A:qos1# show qos sap-ingress 102 detail
```

```

=====
QoS Sap Ingress
=====
-----
Sap Ingress Policy (102)
-----
Policy-id      : 102                      Scope           : Template
Default FC     : be
Criteria-type  : MAC
Sub-Criteria-type : dot1p-only
Accounting     : packet-based
Classifiers Allowed: 32                  Meters Allowed   : 16
Classifiers Used  : 32                  Meters Used     : 16
-----

```

Meter	Mode	CIR Admin CIR Oper	CIR Rule	PIR Admin PIR Oper	PIR Rule	CBS Admin CBS Oper	MBS Admin MBS Oper
1	TrTcm	100 104	closest	200 200	closest	def def	def 500
2	TrTcm	100 104	closest	200 200	closest	def def	def 500
3	TrTcm	100 104	closest	200 200	closest	def def	def 500
4	TrTcm	100 104	closest	200 200	closest	def def	def 500
5	TrTcm	100 104	closest	200 200	closest	def def	def 500
6	TrTcm	100 104	closest	200 200	closest	def def	def 500
7	TrTcm	100 104	closest	200 200	closest	def def	def 500
8	TrTcm	100 104	closest	200 200	closest	def def	def 500
9	TrTcm	100 104	closest	200 200	closest	def def	def 500
10	TrTcm	100 104	closest	200 200	closest	def def	def 500
11	TrTcm	100 104	closest	200 200	closest	def def	def 500
12	TrTcm	100 104	closest	200 200	closest	def def	def 500
13	TrTcm	100 104	closest	200 200	closest	def def	def 500
14	TrTcm	100 104	closest	200 200	closest	def def	def 500
15	TrTcm	100	closest	200	closest	def	def

```

16      TrTcm      104      closest      200      def      500
      100      200      closest      def
      104      200      def      500
-----
FC              UCastM      MCastM      BCastM      UnknownM
-----
be              1              11              16              16
l2              8              16              16              16
af              7              15              16              16
l1              6              14              16              16
h2              5              13              16              16
ef              4              12              16              16
h1              3              10              16              16
nc              2              9              16              16
-----
Match Criteria
-----
Entry              : 1
Src MAC              :
Dst MAC
=====
*A:qos1#

```


Access Egress QoS Policies

In This Section

This section provides information to configure Access Egress QoS policies using the command line interface.

Topics in this section include:

- [Overview on page 262](#)
- [Basic Configurations on page 262](#)
- [Create Access Egress QoS Policies on page 263](#)
- [Default Access Egress QoS Policy Values on page 265](#)

Overview

An access egress policy defines the marking for the traffic egressing on the access ports. Access-egress policies are used at the Ethernet port and define the marking values to use for traffic egressing on the Ethernet port. It defines the marking values to use for a forwarding class (FC).

There is one default access egress policy which is identified as policy ID 1. The default policy can be copied but cannot be deleted or modified. A remarking policy can be defined for each access egress policy and remarking is disabled by default. Only remarking policy of type dot1p, dot1p-lsp-exp-shared, dscp or dot1p-dscp can be used with access-egress policy.

7210 SAS-X supports SAP-based egress marking and port-based egress marking on only access ports. Users have an option to turn on either sap-based marking or port-based marking using the command 'sap-qos-marking' under the CLI configure>port>ethernet>access>egress context. In SAP-based marking the remark policy defined in the SAP egress policy associated with each SAP is used to mark the packets egressing out of SAP if marking is enabled. In port-based marking, the remark policy defined in the access-egress policy associated with the access port determines the marking values to use for all the SAPs defined on that port. SAP-based marking is only supported for L2 SAPs, i.e. SAPs configured in Epipe, VPLS and PBB (I-SAPs only) service. Port-based marking is supported for L3 SAPs (i.e. SAPs configured in VPRN services), PBB B-SAPs and other L2 SAPs. More information on the CLI command 'sap-qos-marking' is available in the 7210 Systems guide. The access egress policy is used only when port-based marking has been enabled (that is, sap-qos-marking is set to disable). By default, sap-based marking is enabled and sap-qos-marking is set to enable.

The default access-egress policy is as shown below

```

-----
*A:Dut-A>config>qos>access-egress# info detail
-----
          description "Default Access egress QoS policy."
          no remarking
          remark 2
-----

```

Basic Configurations

A basic access egress QoS policy must conform to the following:

- Have a unique access egress QoS policy ID.
- Have a description of the policy features.

Create Access Egress QoS Policies

Configuring and applying QoS policies is optional. If no QoS policy is explicitly applied to an access port, a default QoS policy 1 is applied.

Access Egress QoS Policy

To create an access egress policy, you must define the following:

- A new policy ID value. The system will not dynamically assign a value.
- Include a description. The description provides a brief overview of policy features.
- Remark - By default, remarking is disabled.
- A new remarking policy of the appropriate type can be defined for each access egress policy. The policy ID for the remarking policy must be specified.
- Configure port-based marking under the `configure> port>ethernet>access>egress` context.

The following displays the access egress QoS policy configuration:

```
*A:7210-SAS-X>config>qos>access-egress# info detail
-----
      description "policy1"
      remarking
      remark 2
-----
*A:7210-SAS-X>config>qos>access-egress#
```

Applying Access Egress QoS Policies

Apply access egress policies to the following entities:

- Ethernet ports

A policy can be applied to the ports that are in access mode.

Ethernet Ports

Use the following CLI syntax to apply a access-egress policy to an Ethernet port:

CLI Syntax: `config>port#
ethernet access egress
qos access-egress-policy-id
sap-qos-marking disable`

The following output displays the port configuration.

```
*A:7210-SAS-X>config>port>ethernet# info
-----
mode access
  access
    egress
      sap-qos-marking disable
      qos 10
    exit
  exit
-----
*A:7210-SAS-X>config>port>ethernet#
```

Default Access Egress QoS Policy Values

The default access egress policy is identified as policy-id 1. The default policy cannot be edited or deleted. The following displays default policy parameters:

```
*A:7210-SAS-X>config>qos>access-egress# info detail
-----
description "Default Access egress QoS policy."
no remarking
-----
*A:7210-SAS-X>config>qos>access-egress#
```

Editing QoS Policies

Existing policies and entries can be edited through the CLI or NMS. The changes are applied immediately to all services where the policy is applicable.

To prevent configuration errors perform the following:

1. Copy the policy to a work area
2. Edit the policy
3. Over write the original policy

Deleting QoS Policies

Every access Ethernet port is associated, by default, with the default access egress policy (policy-id 1). You can replace the default policy with a customer-configured policy, but you cannot entirely remove the policy from the port configuration. When you remove a non-default access egress policy, the association reverts to the default policy-id 1.

A QoS policy cannot be deleted until it is removed from all access ports where they are applied.

```
*A:7210-SAS-X>config>qos# no access-egress 30  
MINOR: CLI Could not remove Access egress policy "30" because it is in use.
```

Removing a Policy from the QoS Configuration

CLI Syntax: `config>qos# no access-egress policy-id`

Example:

```
config>qos# no access-egress 100  
config>qos# no access-egress 1010
```


Access Egress QoS Policy Command Reference

Command Hierarchies

Configuration Commands

- **config**
- **qos**
 - **access-egress** *policy-id* [**create**]
 - **no access-egress** *policy-id*
 - **description** *description-string*
 - **no description**
 - **remark** *policy-id*
 - **no remark**
 - **marking**
 - **no marking**

Show Commands

- show**
 - **qos**
 - **access-egress** [*policy-id*] [**association**|**detail**]

Configuration Commands

Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>qos>access-egress
Description	This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the context in the configuration file. The no form of this command removes any description string from the context.
Default	No description is associated with the configuration context.
Parameters	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

access-egress

Syntax	access-egress <i>policy-id</i> [create] no access-egress <i>policy-id</i>
Context	config>qos
Description	This command is used to create or edit an access egress QoS policy. The egress policy defines the remark policy for the traffic egressing on the access port. Remarking is disabled by default on the access egress policies. The policy can be applied to multiple access ports. The access egress policy is common to services (SAPs) that are all egressing on a particular port. Any changes made to an existing policy are applied to all access ports on which the policy is specified. 7210 SAS-X supports SAP-based egress marking and port-based egress marking on only access ports. User have an option to turn on either sap-based marking or port-based marking using the command 'sap-qos-marking' under the port context. In SAP-based marking the remark policy

defined in the SAP egress policy associated with each SAP is used to mark the packets egressing out of SAP if marking is enabled. In port-based marking, the remark policy defined in the access-egress policy associated with the access port determines the marking values to use for all the SAPs defined on that port. SAP-based marking is only supported for L2 SAPs, that is, SAPs configured in Epipe, VPLS and PBB (I-SAPs only) service. Port-based marking is supported for both L3 SAPs, that is, SAPs configured in VPRN services, PBB B-SAPs and other L2 SAPs. More information on the CLI command ‘sap-qos-marking’ is available in the 7210 Systems guide. The access egress policy is used only when port-based marking has been enabled (that is, sap-qos-marking is set to disable).

The system uses the access egress policy for marking only if the port with which this policy is associated is enabled for port-based marking (that is, the command sap-qos-marking is set to disable). When port-based marking is enabled, the system is capable of marking all the packets egressing out of the port with either dot1p or dscp or both (that is, both dot1p and dscp). If remarking is enabled and the remark policy is of type ‘dot1p’ or ‘dot1p-lsp-exp-shared’ then the dot1p bits are marked in the packet based on the FC to dot1p values specified in the remark policy. If remarking is enabled and the remark policy is of type ‘dscp’ then the IP DSCP bits are marked in the packet. If remarking is enabled and the remark policy is of type ‘dot1p-dscp’ then both dot1p and IP DSCP bits are marked in the packet.

Note: When port-based marking is enabled and marking for both dot1p and IP DSCP bits is configured, the system marks dot1p and IP DSCP bits for all the packets sent out of both L2 SAPs and L3 SAPs. It is recommended that if both L2 and L3 SAPs are configured on the same port, then remark policy of type dot1p, that marks only dot1p bits be used.

The **no** form of this command deletes the access-egress policy. A policy cannot be deleted until it is removed from all access ports where it is applied. When an access-egress policy is removed from an access port, the access port will revert to the default access-egress policy-id 1.

Parameters *policy-id* — The value that uniquely identifies the access-egress policy.

Values 1 — 65535

create — The keyword used to create an access-egress policy. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

remark

Syntax **remark** *policy-id*
[no] remark

Context config>qos>access-egress

Description This command specifies the remarking policy for the access egress policy.

Only remark policy of type dot1p or dot1p-lsp-exp-shared or dscp or dot1p-dscp is allowed for use with access-egress policy.

Parameters *policy-id* — The value that uniquely identifies the remark policy.

Values 1 — 655353

remarking

Syntax **[no] remarking**
remarking

Context config>qos>access-egress

Description This command enables the system to remark egress packets. The **no** form of the command disables remarking.

If remarking is enabled and no remark policy is explicitly attached then the default remark policy 2 is used.

Default Remarking is disabled by default.

Show Commands

access-egress

- Syntax** **access-egress** [*policy-id*] [**association** | **detail**]
- Context** show>qos
- Description** This command displays Access egress QoS policy information.
- Parameters** *policy-id* — Displays information about the specific policy ID.
- Values** 1 — 65535
- association** — Displays a list of ports on which the policy is applied.
- detail** — Displays detailed policy information including policy associations.
- Access Egress Output** — The following table describes Access egress show command output.

Label	Description
Policy-Id	The ID that uniquely identifies the policy.
Remark	True — Remarking is enabled for the policy. False — Remarking is disabled for the policy.
Description	A text string that helps identify the policy's context in the configuration file
Port-Id	Specifies the physical port identifier that associates the access egress QoS policy.
Remark Pol Id	Displays the policy ID of the remark policy defined for the access egress policy.

Sample Output

```
A:SAS-X>show>qos# access-egress 2

=====
QoS Access Egress
=====
-----
Policy-id      : 2                      Scope      : Template
Remark        : True                   Remark Pol Id: 3
Description    : policy1
```

```

=====
*A:SAS-X>show>qos#

*A:SAS-X>show>qos# access-egress 2 association

=====
QoS Access Egress
=====
-----
Policy-id      : 2                      Scope      : Template
Remark        : True                   Remark Pol Id: 3
Description    : policy1

-----
Associations
-----
Port-id : 1/1/2

=====
*A:SAS-X>show>qos#

*A:SAS-X>show>qos# access-egress 2 detail

=====
QoS Access Egress
=====
-----
Policy-id      : 2                      Scope      : Template
Remark        : True                   Remark Pol Id: 3
Description    : policy1

-----
Associations
-----
Port-id : 1/1/2

=====
*A:SAS-X>show>qos#

```


SAP Egress Policies

In This Section

This section provides information to configure SAP egress policies using the command line interface.

Topics in this section include:

- [Overview on page 278](#)

Overview

The SAP Egress policy defines the Service Level Agreement (SLA) for service packets as they egress on the SAP. Egress SAP QoS policies allow the definition of queue parameters along with remark policy.

By default, the system allocates and associates “8” queues for use by each SAP created in the system. The system predefines the FC to queue map. Each queue must have a CIR rate, PIR rate, and a ‘cir-level’. The pir-level of the queue is assigned by the system based on the cir-level configured by the user. The CIR rate determines the amount of committed bandwidth to be made available to this queue. The PIR rate forces a hard limit on the packets transmitted through the queue. The hardware does not support a linear range of values for the rate parameters (both cir and pir). The user can specify the computation method of rates to match the rates supported by the hardware, through Adaptation-rules.

Each queue is also associated with a scheduler. The scheduler uses the queues cir and pir rate, cir-level, pir-level and pir-weight to distribute the available bandwidth among the queues. When multiple queues are in use, the level parameter determines the scheduling order of the queues. Queues with higher value of level parameter are scheduled out first.

A queue-management policy can be associated with a queue to manage the buffer allocation for in-profile and out-of profile packets. Note: Only a remark policy of type ‘dot1p-only’ or ‘dot1p-lsp-exp-shared’ is available for use with SAP Egress policy.

The following table shows the default FC to queue map assigned by the system.

Table 33: FC to Queue Map

FC	Queue Number
NC	Queue #8
H1	Queue #7
EF	Queue #6
H2	Queue #5
L1	Queue #4
AF	Queue #3
L2	Queue #2
BE	Queue #1

NOTE: For more information on service egress scheduling, refer to the “QoS Schedulers” section.

Configuration Guidelines

- When a SAP is created eight queues are created and associated with it. User cannot create or delete any queues.
 - The ‘pir-level’ of the queue is assigned by the system based on the ‘cir-level’ configured by the user. User cannot change this.
 - The Forwarding Class (FC) to queue map is static and cannot be configured by the user.
-

Basic Configurations

A basic SAP Egress policy must confirm to the following:

- Each SAP Egress must have a unique policy ID.
 - Define the queue parameters (cir, pir, cir-level, pir-weight) for all the queues.
-

Create a SAP Egress Policy

To create a new SAP Egress policy, define the following:

- A SAP Egress policy name.
- Provide a brief description of the policy features.
- Provide the queue parameters for all the queues

Use the following CLI syntax to configure a SAP Egress policy:

CLI Syntax:

```
A:SASX# /configure qos sap-egress
- no sap-egress <policy-id>
- sap-egress <policy-id> [create]
```

```
<policy-id>      : [1..65535]
<create>        : keyword - mandatory while creating an entry.
```

```
[no] description  - Description for this sap-egress policy
                  + Configure a queue
[no] remark       - Specify Remarking policy for this policy
[no] remarking    - Enable/disable remarking
[no] scope        - Specify scope of the policy
```

The following output displays the SAP Egress policy configuration:

```
A:SASX>config>qos>sap-egress# info
-----
description "Default SAP egress QoS policy."
queue 1
exit
queue 2
exit
queue 3
exit
queue 4
exit
queue 5
exit
queue 6
exit
queue 7
exit
queue 8
exit
*A:SAS-X-C>config>qos>sap-egress# info detail
-----
description "Default SAP egress QoS policy."
scope template
no remarking
remark 2
queue 1
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
queue 2
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
queue 3
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
queue 4
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
queue 5
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
queue 6
    port-parent cir-level 1 pir-weight 1
```

```

        adaptation-rule pir closest cir closest
        rate cir 0 pir max
        queue-mgmt "default"
    exit
queue 7
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
queue 8
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
-----
*A:SAS-X-C>config>qos>sap-egress#

A:SASX>config>qos>sap-egress# info detail
-----
description "Default SAP egress QoS policy."
scope template
no remarking
remark 2
queue 1
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
queue 2
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
queue 3
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
queue 4
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"
exit
queue 5
    port-parent cir-level 1 pir-weight 1
    adaptation-rule pir closest cir closest
    rate cir 0 pir max
    queue-mgmt "default"

```

Editing QoS Policies

Existing policies and entries can be edited through the CLI or NMS. The changes are applied immediately to all services where the policy is applicable.

To prevent configuration errors perform the following:

1. Copy the policy to a work area
2. Edit the policy
3. Over write the original policy

Configuration Commands

Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>qos>access-egress
Description	This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the context in the configuration file. The no form of this command removes any description string from the context.
Default	No description is associated with the configuration context.
Parameters	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

SAP Egress Policy Commands

adaptation-rule

Syntax	adaptation-rule [cir <i>adaptation-rule</i>] [pir <i>adaptation-rule</i>] no adaptation-rule
Context	config>qos>sap-egress>queue
Description	This command defines the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint. The no form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific adaptation-rule is removed, the default constraints for cir and pir apply.
Default	adaptation-rule pir closest cir closest
Parameters	<i>adaptation-rule</i> — Specifies the adaptation rule to be used while computing the operational CIR or PIR value. Values pir — Defines the constraints enforced when adapting the PIR rate defined within the queue queue-id rate command. The pir parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the rate command is not specified, the default applies. cir — Defines the constraints enforced when adapting the CIR rate defined within the queue queue-id rate command. The cir parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the cir parameter is not specified, the default constraint applies. max — The max (maximum) option is mutually exclusive with the min and closest options. The hardware step size varies with the configured rate. min — The min (minimum) option is mutually exclusive with the max and closest options. The hardware step size varies with the configured rate. closest — The closest parameter is mutually exclusive with the min and max parameter. The hardware step size varies with the configured rate.

port-parent

Syntax	port-parent [cir-level <i>cir-level</i>] [pir-weight <i>pir-weight</i>] no port-parent
Context	config>qos>sap-egress>queue
Description	The system creates and associates a port-scheduler with every access port on the system. Every queue within a SAP is associated with the port scheduler available on the port on which the SAP is created. This command provides the context to configure the queue parameters ‘cir-level’ and ‘pir-weight’. The port scheduler uses these parameters to apportion the bandwidth to all the queues vying for the available bandwidth. The no form of the command reverts the queue to use the default cir-level and pir-weight values.
Default	Port-parent with default values for cir-level “1” and pir-weight “1”.
Parameters	cir-level <i>cir-level</i> — Specifies the priority of the queue with respect to other queues. The priority of the queue is used only in the CIR loop. Level "8" is the highest priority and level "1" is the lowest priority. Level “8” is treated specially by the schedulers. The scheduler tries to ensure that the configured CIR rate is always made available to queues configured at this level, irrespective of whether CIR for other queues are being met or not. In other words, the system tries to satisfy CIR for all the level-8 queues before it tries to satisfy the CIR of queues configured at other levels. PIR rate configured for level-8 queues is ignored by the system. In the PIR loop, the priority of the queues cannot be configured. The system assigns the priority to the queues based on the configured cir-level. Refer to the QoS scheduler section of the user guide to see the default assignment of pir-levels to the queue corresponding to the cir-level configured by the user. Values 1 — 8 (8 is the highest priority) Default 1 pir-weight <i>pir-weight</i> — Specifies the relative weight of the queue with respect to the other queues. The weight parameter is used only in the PIR loop. If the level parameter of a queue is set to ‘8’, the weight parameter is not used. Values 1 — 100 Default 1

queue

Syntax	queue <i>queue-id</i>
Context	config>qos>sap-egress
Description	This command is used to configure the queue parameters.

Parameters *queue-id* — id of the queue.
Values 1 — 8

queue-mgmt

Syntax **[no] queue-mgmt name**

Context config>qos>sap-egress>queue

Description This command associates the specified queue management policy with this queue.

The queue management policy is used to specify the queue buffer parameters and queue slope policy parameters.

The **no** form of the command associates the default SAP egress queue management policy with this queue.

Parameters name — The name of the queue management policy
Values 1—32 characters

rate

Syntax **rate cir *cir-rate-in-kbps* [pir *pir-rate-in-kbps*]**
no rate

Context config>qos>sap-egress>queue

Description This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the port. Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth. The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth.

The rate command can be executed at anytime, altering the PIR and CIR rates for all queues created on the access ports.

The **no** form of this command returns all queues created with the queue-id by association with the QoS policy to the default PIR and CIR parameters (max, 0).

Default rate cir 0 pir max — The max default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The max value is mutually exclusive to the pir-rate value.

Parameters	cir <i>cir-rate-in-kbps</i> — The cir parameter overrides the default administrative CIR used by the queue. If the rate command is not executed or the cir parameter is not explicitly specified, the default CIR value is used. Values 0 — 10000000, max Default 0 pir <i>pir-rate-in-kbps</i> — Defines the administrative PIR rate, in kilobits, for the queue. When the rate command is executed, a PIR setting is optional. If the rate command is not executed, the default PIR of maximum value is used. Values 1 — 10000000, max Default max
-------------------	--

remark

Syntax	[no] remark <policy-id>
Context	config>qos>sap-egress
Description	This command associates the specified remark policy with this SAP. Remark policies of type dot1p and dot1p-lsp-exp-shared can be specified. It allows the users to specify the dot1p values to use for marking the ethernet header fields of the packets sent out through this SAP. The no form of the command associates the default the remark policy “2” with this SAP.
Default	2
Parameters	<i>policy-id</i> — The ID of the remark policy. Values 1 — 65535

remarking

Syntax	[no] remarking
Context	config>qos>sap-egress
Description	This command is used to enable or disable remarking on service egress. The no form of the command disables remarking on service egress.

sap-egress

Syntax	sap-egress <policy-id> [create]
Context	config>qos
Description	This command enables the context to configure a SAP Egress policy. The SAP egress policy determines the QoS treatment to packets at service egress. The system creates and associates eight queues to each of the SAPs in the system. User cannot create or delete a queue. SAP egress policy allows the user to define the queue parameters for the eight queues.
Default	1
Parameters	<i>policy-id</i> — The ID of the SAP Egress policy Values 1 — 65535

scope

Syntax	scope {exclusive template} no scope
Context	config>qos>sap-egress
Description	This command configures the scope as exclusive or template. The no form of this command sets the scope of the policy to the default of template.
Default	template
Parameters	exclusive — When the scope of a policy is defined as exclusive, the policy can only be applied to one interface. If a policy with an exclusive scope is assigned to a second interface an error message is generated. If the policy is removed from the exclusive interface, it will become available for assignment to another exclusive interface. template — When the scope of a policy is defined as template, the policy can be applied to multiple interfaces on the router. Default QoS policies are configured with template scope. An error is generated if you try to modify the scope parameter from template to exclusive scope on default policies.

Operational Commands

copy

Syntax	copy sap-egress <src-pol> <dst-pol> [overwrite]
Context	config>qos
Description	This command copies the existing SAP egress QoS policy entries to another SAP egress QoS policy. The copy command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the overwrite keyword. If the destination policy already exists, the key word overwrite must be specified.
Parameters	<i>src-pol</i> — Specifies the source policy. Values 1—65535 <i>dst-pol</i> — Specifies the destination policy. Values 1—65535 overwrite — The information in the destination policy is overwritten by the information in the source policy.

Show Commands

sap-egress

- Syntax** **sap-egress** [*policy-id*] [**association** | **detail**]
- Context** show>qos
- Description** This command displays sap egress QoS policy information.
- Parameters** *policy-id* — Displays the policy id of the sap-egress policy.
association — Displays associations related to the specified sap-egress policy.
detail — Displays detailed policy information including the policy associations.
SAP Egress Output — The following table describes Access egress show command output.

Label	Description
Policy-Id	The ID that uniquely identifies the policy.
Remark	True — Remarking is enabled for all the Dot1q-tagged packets that egress the ports on which the sap-egress QoS policy is applied and remarking is enabled. False — Remarking is disabled for the policy.
Remark Pol Id	Displays the policy id of the remarking policy.
Accounting	Specifies whether the accounting mode is packet-based or frame-based.
Scope	Exclusive — Implies that this policy can be applied only to a single access egress port. Template — Implies that this policy can be applied to multiple access ports on the router.
Description	Template — Implies that this policy can be applied to multiple access ports on the router. A text string that helps identify the policy’s context in the configuration file

Queue Rates and Rules

Label	Description (Continued)
QueueId	Displays the Queue identifier associated with the sap-egress QoS policy.
Explicit/Default	Explicit — Specifies the egress IEEE 802.1P (dot1p) bits marking for fc-name if explicitly configured. Default — Specifies the default dot1p value according to FC-Dot1p marking map as defined in Table 17, Default SAP Egress Policy ID 1 Definition, on page 48 if explicit values are not configured..
CIR	Specifies the administrative Committed Information Rate (CIR) parameters for the queue. The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth.
CIR Adpt Rule	min — The operational CIR for the queue will be equal to or greater than the administrative rate specified using the rate command. max — The operational CIR for the queue will be equal to or less than the administrative rate specified using the rate command. closest — The operational CIR for the queue will be the rate closest to the rate specified using the rate command without exceeding the operational PIR.
PIR	Specifies the administrative Peak Information Rate (PIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the access port.
PIR Adpt Rule	min — The operational PIR for the queue will be equal to or greater than the administrative rate specified using the rate command. max — The operational PIR for the queue will be equal to or less than the administrative rate specified using the rate command. closest — The operational PIR for the queue will be the rate closest to the rate specified using the rate command.

Parent Details

QueueId	Displays the Queue identifier associated with the sap-egress QoS policy
---------	---

Label	Description (Continued)
Port	Indicates if the parent scheduler is port scheduler or not.
CIR Level	Displays the priority of the queue in the CIR loop.
PIR Weight	Displays the weight of the queue used in the PIR loop.

High Slope/Low slope

QueueId	Displays the Queue identifier associated with the sap-egress QoS policy.
State	Displays the state of the queue. The state of the queue can be either “Up” or “Down”
Start Avg	Specifies the low priority or high priority RED slope position for the shared buffer average utilization value where the packet discard probability starts to increase above zero.
Max Avg	Specifies the percentage of the shared buffer space for the buffer pool at which point the drop probability becomes “1”. This parameter is expressed as a decimal integer.
Max Prob	Specifies the high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one.

Burst Sizes and Time Average Factor

QueueId	Displays the Queue identifier associated with the sap-egress QoS policy
CBS	Displays the configured CBS value
MBS	Displays the configured MBS value
Time Average Factor	Displays the value of the time average factor in use
Queue-Mgmt	Displays the Queue management policy in use

Service Associations

Service-Id	The unique service ID number which identifies the service in the service domain.
------------	--

Label	Description (Continued)
Customer - Id	Specifies the customer ID which identifies the customer to the service.
SAP	Specifies the Service Access Point (SAP) within the Service where the SAP egress policy is applied.

Sample Output

```
A:SASX>config>qos# show qos sap-egress 1 detail
```

```
=====
QoS Sap Egress
=====
-----
Sap Egress Policy (1)
-----
Scope                : Template
Remark               : False          Remark Pol Id        : 2
Accounting           : frame-based
Description          : Default SAP egress QoS policy.
-----
Queue Rates and Rules
-----
-----
QueueId      CIR      CIR Adpt Rule      PIR      PIR Adpt Rule
-----
Queue1       0        closest           max      closest
Queue2       0        closest           max      closest
Queue3       0        closest           max      closest
Queue4       0        closest           max      closest
Queue5       0        closest           max      closest
Queue6       0        closest           max      closest
Queue7       0        closest           max      closest
Queue8       0        closest           max      closest
-----
Parent Details
-----
-----
QueueId      Port      CIR Level      PIR Weight
-----
Queue1       True      1              1
```

Queue2	True	1	1
Queue3	True	1	1
Queue4	True	1	1
Queue5	True	1	1
Queue6	True	1	1
Queue7	True	1	1
Queue8	True	1	1

High Slope

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	70	90	75
Queue2	Down	70	90	75
Queue3	Down	70	90	75
Queue4	Down	70	90	75
Queue5	Down	70	90	75
Queue6	Down	70	90	75
Queue7	Down	70	90	75
Queue8	Down	70	90	75

Low Slope

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Down	50	75	75
Queue2	Down	50	75	75
Queue3	Down	50	75	75
Queue4	Down	50	75	75
Queue5	Down	50	75	75
Queue6	Down	50	75	75
Queue7	Down	50	75	75
Queue8	Down	50	75	75

Burst Sizes and Time Average Factor

QueueId	CBS	MBS	Time Average Factor	Queue-Mgmt
Queue1	def	def	7	default

Queue2	def	def	7	default
Queue3	def	def	7	default
Queue4	def	def	7	default
Queue5	def	def	7	default
Queue6	def	def	7	default
Queue7	def	def	7	default
Queue8	def	def	7	default

Associations

Service-Id : 1 (Epipe) Customer-Id : 1
- SAP : 1/1/1:1
Service-Id : 101 (Epipe) Customer-Id : 1
- SAP : 1/1/2:101

Mirror SAPs

No Mirror SAPs Found.

=====
A:SASX>config>qos#

SAP Egress Policy Command Reference

Command Hierarchies

Configuration Commands

```
show
  — qos
    — sap-egress <policy-id> create
      — [no] description description-string
      — queue <queue-id>
        — [no] port-parent [cir-level level] [pir-weight weight]
        — [no] adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
        — [no] rate cir <rate> pir<rate>
        — queue-mgmt <name>
        — no queue-mgmt
      — scope {exclusive | template}
      — [no] remark <policy-id>
      — [no] remarking
```

Copy Commands

```
show
  — qos
    — copy sap-egress <src-pol> <dst-pol> [overwrite]
```

Show Commands

```
show
  — qos
    — sap-egress [<policy-id>] [detail | association]
```


QoS Schedulers

In This Section

This section provides information about the scheduler support available in the 7210 SAS X.

Topics in this section include:

- [Overview on page 300](#)

Overview

The system creates a port scheduler for all the ports (both access ports and network ports) in the system and distributes the available bandwidth to all the queues using that port to send out packets. The port scheduler works either at line-rate or configured egress rate. The port scheduler allocates bandwidth to the SAPs configured on the access port or to the queues configured on the network port.

The system creates a SAP scheduler for all the SAPs in the system. It distributes the available port bandwidth to all the queues allocated to a SAP. The SAP scheduler allocates bandwidth to the queues based on the configured CIR, PIR rates and the priority and weight assigned to the queues.

Note: The SAP aggregate rate can be configured for the SAP scheduler when SAP based shaping is in use.

Listed below are the two passes made by the port scheduler:

- belowCIR pass
- aboveCIR pass

In the belowCIR pass, the port scheduler distributes the available port bandwidth among all the queues created on the port based on the configured cir-level, until the configured CIR rate is met. Queues with higher cir-level are scheduled first. The system services queues with CIR-Level “8” prior to servicing other queues in the system.

Note: To ensure that all the queues in the system are serviced by the scheduler, the queues with cir-level “8” must be capped to a pre-determinate rate.

In the aboveCIR pass, the port scheduler distributes the remaining port bandwidth among all the queues based on its pir-level and pir-weight.

The pir-level assigned by the system is given in [Table 34](#):

Table 34: pir-level Assignments

CIR Level Config- ured for use in CIR loop	PIR level assigned by system for use in PIR loop
8(highest)	N/A
7	4
6	3
5	3

Table 34: pir-level Assignments (Continued)

CIR Level Config- ured for use in CIR loop	PIR level assigned by system for use in PIR loop
4	2
3	2
2	1
1(lowest)	1(lowest)

Queue Management Policies

In This Section

This section provides information to configure queue management policies using the command line interface.

Topics in this section include:

- [Overview on page 304](#)
- [Basic Configurations on page 305](#)

Overview

A set of profiles or templates are available in hardware for configuring the queue parameters such as CBS, MBS, and WRED slopes parameters per queue. These profiles are available for use with multiple queues of the system. Queue management policy allows the user to define the queue parameters and allow sharing among the queues.

A single system buffer pool is available for use by all the queues in the system. Users can allocate the amount of buffers that each queue can use by specifying the CBS and MBS parameters in the queue management policy.

Weighted Random Early Detection (WRED) is available to manage buffers during periods of congestion. WRED slopes are supported for each queue in the system. The Queue Management Policies allow the user to configure slope parameters that dictate a WRED profile for each queue. Each queue supports two slopes:

- Slope for in-profile or high priority traffic.
- Slope for out-of-profile or low priority traffic.

Each slope allows specifying the start-average, the max-average, the drop-probability and the Time Average Factor (TAF). Each queue has a default slope policy. Multiple queues in the system can share a single policy. If a policy is shared the system computes the WRED drop probabilities for each of the queues separately based on their average queue length.

Basic Configurations

A basic queue management policy must confirm to the following:

- Each slope policy must have a unique policy ID.
 - High slope and low slope are shut down by default.
 - Default values can be modified but parameters cannot be deleted.
-

Creating a Queue Management Policy

To create a new queue management policy, define the following:

- A queue management policy name.
- Provide a brief description of the policy features.
- Provide CBS and MBS values for default queue management policy.
- The high slope for the high priority WRED slope graph.
- The low slope for the low priority WRED slope graph.
- The time average factor (TAF).
- Slope parameters such as max-avg, start-avg, max-prob, time-average-factor have default values.

Use the following CLI syntax to configure a queue management policy:

CLI Syntax: `config>qos`

```

queue-mgmt name
description description-string
cbs kbytes
mbs kbytes
high-slope
    start-avg percent
    max-avg percent
    max-prob percent
    no shutdown
low-slope
    start-avg percent
    max-avg percent
    max-prob percent
    no shutdown
time-average-factor taf
  
```

The following output displays the queue management policy configuration:

```
A:7210-x>config>qos>queue-mgmt# info
```

```
-----
      high-slope
        shutdown
        start-avg 40
        max-avg 50
      exit
      low-slope
        shutdown
        start-avg 40
        max-avg 80
      exit
    cbs 5000
    mbs 800000
time-average-factor 7
-----
```

Editing QoS Policies

Existing policies and entries can be edited through the CLI or NMS. The changes are applied immediately to all services where the policy is applicable.

To prevent configuration errors perform the following:

1. Copy the policy to a work area
2. Edit the policy
3. Over write the original policy

Queue Management Policy Command Reference

Command Hierarchies

Configuration Commands

```
config
  — qos
    — [no] queue-mgmt name
      — description description-string
      — no description
      — [no] cbs kbytes
      — [no] mbs kbytes
      — [no] high-slope
        — [no] max-avg percent
        — [no] max-prob percent
        — [no] shutdown
        — [no] start-avg percent
      — [no] low-slope
        — [no] max-avg percent
        — [no] max-prob percent
        — [no] shutdown
        — [no] start-avg percent
      — [no] time-average-factor value
```

Show Commands

```
show
  — qos
    — queue-mgmt [<name>] [detail]
```

Operational Commands

```
config
  — qos
    — copy queue-mgmt <src-name> <dst-name> [overwrite]
```

Configuration Commands

Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>qos>queue-mgmt
Description	This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the context in the configuration file. The no form of this command removes any description string from the context.
Default	No description is associated with the configuration context.
Parameters	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands

copy

Syntax	copy queue-mgmt <src-name> <dst-name> [overwrite]
Context	config>qos
Description	This command copies the existing Queue management policy entries to another Queue management policy. The copy command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the overwrite keyword.
Parameters	<i>src-name</i> — Specifies the name of the source policy. Values 1 — 32 characters <i>dst-name</i> — Specifies the name of the destination policy. Values 1 — 32 characters overwrite — The information in the destination policy is overwritten by the information in the source policy.

Queue Management Policy QoS Commands

cbs

Syntax	[no] cbs <i>kbytes</i>
Context	config>qos>queue-mgmt
Description	This command specifies the CBS value (Minimum depth of the queue in kilo bytes).
Parameters	<i>kbytes</i> — Specifies the minimum depth of the queue in kilo bytes.
Values	0 — 500000 default

high-slope

Syntax	[no] high-slope
Context	config>qos>queue-mgmt
Description	This command is used to configure the in-profile WRED slope parameters.

low-slope

Syntax	[no] low-slope
Context	config>qos>queue-mgmt
Description	This command is used to configure the out-of-profile WRED slope parameters.

mbs

Syntax	[no] mbs <i>kbytes</i>
Context	config>qos>queue-mgmt
Description	This command specifies the MBS value (Maximum depth of the queue in kilo bytes).
Parameters	<i>kbytes</i> — Specifies the minimum depth of the queue in kilo bytes.
Values	1 — 500000 default

queue-mgmt

Syntax	[no] queue-mgmt <i>name</i>
Context	config>qos
Description	This command enables the context to configure a QoS queue management policy. A set of profiles/templates are available in hardware for configuring the queue parameters such as CBS, MBS, and WRED slopes parameters per queue. These profiles are available for use with multiple queues of the system. Queue management policy allows the user to define the queue parameters and allow for sharing among the queues.
Parameters	<i>name</i> — The name of the queue management policy. Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

time-average-factor

Syntax	time-average-factor <i>value</i> no time-average-factor
Context	config>qos>queue-mgmt
Description	This command is used to configure the time-average factor.
Default	7
Parameters	<i>value</i> — Represents the Time Average Factor (TAF), expressed as a decimal integer. The value specified for TAF affects the speed at which the shared buffer average utilization tracks the instantaneous shared buffer utilization. A low value weights the new shared buffer average utilization calculation more to the shared buffer instantaneous utilization, zero using it exclusively. A high value weights the new shared buffer average utilization calculation more to the previous shared buffer average utilization value. Values 0 — 15

WRED Slope Commands

max-avg

Syntax	max-avg <i>percent</i> no max-avg
Context	config>qos>queue-mgmt>high-slope config>qos>queue-mgmt>low-slope
Description	This command is used to configure the maximum average value.
Default	max-avg 90 — High slope default is 90% buffer utilization. max-avg 75 — Low slope default is 75% buffer utilization.
Parameters	<i>percent</i> — Specifies the maximum average for the high or low slopes. Values 0 — 100

max-prob

Syntax	max-prob <i>percent</i> no max-prob
Context	config>qos>queue-mgmt>high-slope config>qos>queue-mgmt>low-slope
Description	This command is used to configure the maximum probability value.
Default	max-avg 75 — High slope default is 75% maxium drop probability. max-avg 75 — Low slope default is 75% maxium drop probability.
Parameters	<i>percent</i> — Specifies the maximum probability for the high or low slopes. Values 1 — 99

shutdown

Syntax	[no] shutdown
Context	config>qos>queue-mgmt>high-slope config>qos>queue-mgmt>low-slope
Description	This command enables or disables the administrative status of the WRED slope.

By default, all slopes are shutdown and have to be explicitly enabled (**no shutdown**).

The **no** form of this command administratively enables the WRED slope.

Default **shutdown** - WRED slope disabled implying a zero (0) drop probability

start-avg

Syntax **start-avg** *percent*
 no start-avg

Context config>qos>queue-mgmt>high-slope
 config>qos>queue-mgmt>low-slope

Description This command is used to configure the starting average value.

Default **max-avg 70** — High slope default is 70% buffer utilization.
 max-avg 50 — Low slope default is 50% buffer utilization.

Parameters *percent* — Specifies the starting average for the high or low slopes.
 Values 0 — 100

Show Commands

queue-mgmt

Syntax `queue-mgmt [<name>] [detail]`

Context `show>qos`

Description This command displays queue management policy information.

Parameters *name* — The name of the queue management policy.

detail — Displays detailed information about the queue management policy.

Output **Queue Management Policy Output Fields** — The following table describes queue management policy output fields.

Table 35: Show Queue Management Policy Output Fields

Label	Description
Policy	The ID that uniquely identifies the policy.
Description	A string that identifies the policy's context in the configuration file.
Time Avg	The weighting between the previous shared buffer average utilization result and the new shared buffer utilization.
CBS	Displays the committed burst size.
MBS	Displays the maximum burst size.
Slope Parameters	
Start Avg	Specifies the low priority or high priority RED slope position for the shared buffer average utilization value where the packet discard probability starts to increase above zero.
Max Avg	Specifies the percentage of the shared buffer space for the buffer pool at which point the drop probability becomes 1, expressed as a decimal integer
Admin State	Up — The administrative status of the RED slope is enabled. Down — The administrative status of the RED slope is disabled. Specifies the low priority or high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one.

Table 35: Show Queue Management Policy Output Fields (Continued)

Label	Description
Max Prob.	Specifies the high priority RED slope position for the maximum non-one packet discard probability value before the packet discard probability rises directly to one.
Service Associations	
SAP Egress Policy Id	Displays the SAP Egress policy Id.
Queue Ids	Displays the Queue Ids

Sample Output

```
*A:SASX>config>qos>queue-mgmt# show qos queue-mgmt 20 detail
```

```
=====
QoS Queue Management Policy
=====
Policy          : 20
Description     : (Not Specified)
CBS             : 1                               MBS             : 10
Time Avg       : 7

-----
High Slope Parameters
-----
Start Avg       : 70                               Admin State    : Disabled
Max Avg        : 90                               Max Prob.     : 75

-----
Low Slope Parameters
-----
Start Avg       : 50                               Admin State    : Disabled
Max Avg        : 75                               Max Prob.     : 75

-----
Associations
-----
SAP Egress
-----
SAP Egress Policy Id      : 200
Queue Ids                 : 1, 2, 3, 4, 5, 6, 7, 8
-----

-----
Associations
-----
No Associations Found.
-----
```

Network Queues

No Network Queue Policy Associations found.

Remark Policies

In This Section

This section provides information to configure remark policies using the command line interface.

Topics in this section include:

- [Overview on page 320](#)
- [Basic Configurations on page 322](#)

Overview

The remark policies are used to configure the marking behavior for the system at the egress of access SAPs, network ports and network IP interfaces. These policies allow the user to define the forwarding class to egress marking values and allow them to use the available hardware resources efficiently. Based on the packet encapsulation used, the remark policy allows the user to define and associate appropriate policies to service egress and network egress QoS policies. The 7210 SAS supports the use of the following types of remark policies for different QoS policies:

1. dot1p - Used for service egress and network qos [port type] policies.
2. dscp - Used for network qos [port type] policies.
3. lsp-exp - Used for network qos [ip-interface type] policies.
4. dot1p-dscp - Used for network qos [port type] policies.
5. dot1p-lsp-exp-shared - Used for service egress and network qos [ip-interface type] policies.

The 7210 SAS uses a common pool of hardware resources to mark dot1p values for service egress and lsp-exp values for network IP interface. Remark policies of type dot1p-lsp-exp-shared allow the user to specify one remark policy, for which the system allocates one entry from the common pool, and this entry is used at both service egress and network IP interface.

The type of the remark policy identifies the bits marked in the packet header. Each of these remark policy types can be associated with only appropriate QoS policies and service entities as listed in [Table 36](#).

Table 36: Summary of remark policy and attachment points

Remark Policy Type	Qos Policy	Attachemtn Point	Packet Header Bits Marked
dot1p	SAP-egress policy, Network policy (port)	SAP, Network Port	• Dot1p bits in the L2 header for service packets sent out of a SAP • Dot1p bits in the L2 header for IP and MPLS packets sent out of network port
dscp	Network policy (port)	Network Port	DSCP bits in the IP header for IP packets sent out of network port
lsp-exp	Network policy (IP interface)	Network IP interface	EXP bits in the MPLS header for MPLS packets sent out of network port

Table 36: Summary of remark policy and attachment points

Remark Policy Type	Qos Policy	Attachemtn Point	Packet Header Bits Marked
dot1p-lsp-exp-shared	SAP-egress policy, Network policy (IP interface)	SAP,Network IP Interface	• Dot1p bits in the L2 header for service packets sent out of SAP • EXP bits in the MPLS header for MPLS packets sent out of network port.

Basic Configurations

A basic remark policy must confirm to the following:

- Each remark policy must have a unique policy ID.
 - The remark policy type must be specified.
 - The forwarding class to egress marking values must be specified.
-

Creating a Remark Policy

To create a new remark policy, define the following:

- A remark policy name and type.
- Provide a brief description of the policy features.
- Specify the forwarding class to egress marking values.

Use the following CLI syntax to configure a remark policy:

CLI Syntax:

```
A:7210-x>config>qos# remark 122 remark-type dscp create
```

The following output displays the remark policy configuration:

```
A:7210-x>config>qos>remark# info
-----
fc af
  dscp-in-profile cs1
  dscp-out-profile cp3
exit
fc be
  dscp-in-profile nc2
  dscp-out-profile af11
exit
fc ef
  dscp-in-profile cp1
  dscp-out-profile cp2
exit
fc h1
  dscp-in-profile cp9
  dscp-out-profile cp4
exit
fc h2
  dscp-in-profile cp1
```

```
    dscp-out-profile cp2
exit
fc l1
    dscp-in-profile cp9
    dscp-out-profile cp4
exit
fc l2
    dscp-in-profile cp9
    dscp-out-profile cp4
exit
fc nc
    dscp-out-profile af11
exit
-----
A:7210-x>config>qos>remark#
```

Editing QoS Policies

Existing policies and entries can be edited through the CLI or NMS. The changes are applied immediately to all services where the policy is applicable.

To prevent configuration errors perform the following:

1. Copy the policy to a work area
2. Edit the policy
3. Over write the original policy

Remark Policy Command Reference

Command Hierarchies

Configuration Commands

```
config
  — qos
    — [no] remark <policy-id> remark-type {dot1p | dscp | lsp-exp| dot1p- lsp-exp-shared | dot1p-dscp}
      create
    — [no] description description-string
    — [no] fc fc-name
      — dot1p-in-profile dot1p-value
      — no dot1p-in-profile
      — dot1p-out-profile dot1p-value
      — no dot1p-out-profile
      — dscp-in-profile dscp-value
      — no dscp-in-profile
      — dscp-out-profile dscp-value
      — no dscp-out-profile
      — lsp-exp-in-profile lsp-exp-value
      — no lsp-exp-in-profile
      — lsp-exp-out-profile lsp-exp-value
      — no lsp-exp-out-profile
      — dot1p-lsp-exp-shared-in-profile dot1p-lsp-exp-value
      — no dot1p-lsp-exp-shared-in-profile
      — dot1p-lsp-exp-shared-out-profile dot1p-lsp-exp-value
      — no dot1p-lsp-exp-shared-out-profile
```

Show Commands

```
show
  — qos
    — remark-policy [<policy-id>] [association|detail]
```

Operational Commands

```
config
  — qos
    — copy remark <src-pol> <dst-pol> [overwrite]
```

Configuration Commands

Generic Commands

description

Syntax **[no] description** *description-string*

Context config>qos>remark

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Default No description is associated with the configuration context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Operational Commands

copy

Syntax **copy remark** <*src-pol*> <*dst-pol*> [**overwrite**]

Context config>qos

Description This command copies existing remark policy entries to another remark policy.

The **copy** command is a configuration level maintenance tool used to create new policies using existing policies. It also allows bulk modifications to an existing policy with the use of the **overwrite** keyword.

If the destination policy already exists, the key word **overwrite** must be specified.

Parameters *src-pol* — Specifies the source policy.

Values 1—65535

dst-pol — Specifies the destination policy.

Values 1—65535

overwrite — The information in the destination policy is overwritten by the information in the source policy.

Remark Policy QoS Commands

remark

Syntax	[no] remark <policy-id> remark-type {dot1p dscp lsp-exp dot1p- lsp-exp-shared dot1p-dscp} create
Context	config>qos
Description	This command creates a new remark policy of the specified type. The 7210 SAS allows the sharing of this policy between different service entities to optimize the hardware resources used.

The following types of remark policies are available:

- dot1p (Used for service egress and network qos [port type] policies)
- dscp (Used for network qos [port type] policies)
- lsp-exp (Used for network qos [ip-interface type] policies)
- dot1p-lsp-exp-shared (Used for service egress and network qos [ip-interface type] policies)
- dot1p-dscp (Used for network qos [port type] policies).

The device uses a common pool of hardware resources to mark dot1p values for service egress and lsp-exp values for network IP interface. Remark policies of type dot1p-lsp-exp-shared allow the user to define a single policy, the system allocates a single hardware resource and this policy can be used for multiple SAPs or network IP interfaces. The users can define a single policy of type dot1p or lsp-exp and use it with multiple SAPs and network IP interfaces respectively.

The ‘remark-type’ of the policy also determines the values user is allowed to configure in the policy. For example, if remark-type is dot1p, user is allowed to only specify the forwarding class to dot1p values.

Default	dot1p-lsp-exp-shared
Parameters	policy-id — The policy ID of the remark policy.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

remark-type — Specifies the type of marking values in the remark policy.

Values dot1p — Specifies FC to 802.1 Dot1p value.
lsp-exp— Specifies FC to MPLS LSP EXP values
dot1p-lsp-exp-shared— Specifies FC to both 802.1 Dot1p and MPLS LSP EXP values.

dscp— Specifies the FC to IP DSCP values
dot1p-dscp - Specifies the FC to both 802.1 Dot1p and IP DSCP values.

fc

Syntax [no] fc *fc-name*

Context config>qos>remark

Description This command specifies the forwarding class name and provides the context to configure the marking value for the FC. Based on the type of remark policy created, the FC command allows the user to specify the appropriate marking values. The fc command overrides the default parameters for the forwarding class to the values defined.

The no form of the command removes the forwarding class lsp-exp/dot1p/dscp/dot1p-LSP-EXP map associated with the fc. The forwarding class reverts to the defined parameters in the default remark policy.

Default none

Parameters *fc-name* — Specifies a case-sensitive system-defined forwarding class name for which policy entries are created.

Values be, l2, af, l1, h2, ef, h1, nc

Remark Policy Forwarding Class Commands

dot1p-in-profile

Syntax	dot1p-in-profile <i>dot1p-value</i> no dot1p-in-profile
Context	config>qos>remark>fc
Description	This command specifies the dot1p in-profile value.
Parameters	<i>dot1p-value</i> — A 3-bit value expressed as a decimal integer. This value is used for the IEEE 802.1 dot1p bits in the vlan tag of the ethernet header. Values 0 — 7

dot1p-out-profile

Syntax	dot1p-out-profile <i>dot1p-value</i> no dot1p-out-profile
Context	config>qos>remark>fc
Description	This command specifies the dot1p out-of-profile value.
Parameters	<i>dot1p-value</i> — A 3-bit value expressed as a decimal integer. This value is used for the IEEE 802.1 dot1p bits in the vlan tag of the ethernet header. Values 0 — 7

dscp-in-profile

Syntax	dscp-in-profile <i>dscp-value</i> no dscp-in-profile
Context	config>qos>remark>fc
Description	This command specifies the dscp in-profile value.
Parameters	<i>dscp-value</i> — A 5-bit value expressed as a decimal integer. This value is used for the IP DSCP bits in the IP header. Values 0 — 63

dscp-out-profile

Syntax	dscp-out-profile <i>dscp-value</i> no dscp-out-profile
Context	config>qos>remark>fc
Description	This command specifies the dscp out-profile value.
Parameters	<i>dscp-value</i> — A 5-bit value expressed as a decimal integer. This value is used for the IP DSCP bits in the IP header. Values 0 — 63

lsp-exp-in-profile

Syntax	lsp-exp-in-profile <i>lsp-exp-value</i> no lsp-exp-in-profile
Context	config>qos>remark>fc
Description	This command specifies the in-profile LSP EXP value for the forwarding class. This value is used for all in-profile LSP labeled packets which require marking the egress on the forwarding class queue. When multiple LSP EXP values are associated with the forwarding class at network egress, the last name entered overwrites the previous value. The no form of the command reverts to the factory default in-profile EXP setting.
Parameters	<i>lsp-exp-value</i> — A 3-bit LSP EXP bit value expressed as a decimal integer. Values 0 — 7

lsp-exp-out-profile

Syntax	lsp-exp-out-profile <i>lsp-exp-value</i> no lsp-exp-out-profile
Context	config>qos>remark>fc
Description	This command specifies the in-profile LSP EXP value for the forwarding class. This value is used for all out-of-profile LSP labeled packets which require marking the egress on the forwarding class queue. When multiple LSP EXP values are associated with the forwarding class at network egress, the last name entered overwrites the previous value.

The **no** form of the command reverts to the factory default in-profile EXP setting.

Parameters *lsp-exp-value* — A 3-bit LSP EXP bit value expressed as a decimal integer.

Values 0 — 7

dot1p-lsp-exp-shared-in-profile

Syntax **dot1p-lsp-exp-in-profile** *dot1p-lsp-exp-value*
no dot1p-lsp-exp-in-profile

Context config>qos>remark>fc

Description This command specifies the in-profile Dot1p LSP EXP value for the forwarding class. This value is used for all in-profile LSP labeled packets which require marking the egress on the forwarding class queue.

When multiple Dot1p LSP EXP values are associated with the forwarding class at network egress, the last name entered overwrites the previous value.

The **no** form of the command reverts to the factory default in-profile Dot1p LSP EXP setting.

Parameters *dot1p-lsp-exp-value* — A 3-bit Dot1p LSP EXP bit value, expressed as a decimal integer.

Values 0 — 7

dot1p-lsp-exp-shared-out-profile

Syntax **dot1p-lsp-exp-out-profile** *dot1p-lsp-exp-value*
no dot1p-lsp-exp-out-profile

Context config>qos>remark>fc

Description This command specifies the in-profile Dot1p LSP EXP value for the forwarding class. This value is used for all out-of-profile LSP labeled packets which require marking the egress on the forwarding class queue.

When multiple Dot1p LSP EXP values are associated with the forwarding class at network egress, the last name entered overwrites the previous value.

The **no** form of the command reverts to the factory default in-profile Dot1p LSP EXP setting.

Parameters *dot1p-lsp-exp-value* — A 3-bit Dot1p LSP EXP bit value, expressed as a decimal integer.

Values 0 — 7

Show Commands

remark-policy

Syntax **remark-policy** [*<policy-id>*] [**association**|**detail**]

Context show>qos

Description This command displays remark policy information.

Parameters *policy-id* — The ID of the remark policy.

detail — Displays detailed information about the remark policy.

Output **Remark Policy Output Fields** — The following table describes remark policy output fields.

Table 37: Show Remark Policy Output Fields

Label	Description
Policy ID	The ID that uniquely identifies the policy.
Remark Policy-id	Displays the policy-id of the remark policy.
Type	Displays the type of remark policy.
Description	A string that identifies the policy's context in the configuration file.
FC Name	Specifies the forwarding class name.
dot1P/LSP EXP In value	dot1p/LSP EXP value for in-profile packets.
dot1P/LSP EXP Out value	dot1p/LSP EXP value for out-of-profile packets.
DCSP In value	DSCP value used for in-profile packets
DCSP Out value	DSCP value used for out-of-profile packets
Service Associations	
SAP Egress Policy Id	Displays the policy ID of the SAP Egress policy.
Service-Id	The unique service ID number which identifies the service in the service domain.
Customer-Id	Specifies the customer ID which identifies the customer to the service.

Table 37: Show Remark Policy Output Fields (Continued)

Label	Description
SAP	Specifies the a Service Access Point (SAP) within the service where the SAP ingress policy is applied.
Network	
Network Policy Id	Displays the network policy Id.
Interface Association	
Interface	Displays the associated interface.
IP Addr.	Displays the IP address of the interface.

Sample Output

```
*A:SAS-X-C>config>qos# show qos remark-policy
```

```
=====
SAS Remarking Policies
=====
Policy-Id Type Description
-----
1 dscp Default Remarking Policy for DSCP
2 dot1p-lsp-exp-shared Default Remarking Policy for dot1P and LSP*
500 dot1p-lsp-exp-shared
505 dot1p-lsp-exp-shared
510 dot1p-lsp-exp-shared
515 dot1p-lsp-exp-shared
520 dot1p-lsp-exp-shared
525 dot1p-lsp-exp-shared
530 dot1p-lsp-exp-shared
535 dot1p-lsp-exp-shared
540 dot1p-lsp-exp-shared
545 dot1p-lsp-exp-shared
550 dot1p
555 dot1p
560 dot1p
565 dot1p
570 dot1p
=====
```

- indicates that the corresponding row element may have been truncated.

```
*A:SAS-X-C>config>qos# show qos remark-policy 500 association
```

```
=====
QoS Remarking Policies
=====
Remark Policy-id : 500 Type : dot1p-lsp-exp-shared
Description : (Not Specified)
```

```

-----
Associations
-----
SAP Egress
-----
SAP Egress Policy Id          : 5001
-----

Associations
-----
Service-Id      : 500 (VPLS)          Customer-Id : 1
- SAP : lag-2:500

SAP Egress Policy Id          : 5701
-----

Associations
-----
Service-Id      : 570 (VPLS)          Customer-Id : 1
- SAP : lag-2:570

SAP Egress Policy Id          : 6401
-----

Associations
-----
Service-Id      : 640 (VPLS)          Customer-Id : 1
- SAP : lag-2:640

SAP Egress Policy Id          : 10001
-----

Associations
-----
Service-Id      : 1000 (VPLS)         Customer-Id : 1
- SAP : lag-2:1000

SAP Egress Policy Id          : 10701
-----

Associations
-----
Service-Id      : 1070 (VPLS)         Customer-Id : 1
- SAP : lag-2:1070

SAP Egress Policy Id          : 11401
-----

Associations
-----
Service-Id      : 1140 (VPLS)         Customer-Id : 1
- SAP : lag-2:1140

SAP Egress Policy Id          : 15001
-----

Associations
-----

```

Service-Id : 1500 (VPLS) Customer-Id : 1
- SAP : lag-4:1500

SAP Egress Policy Id : 15701

Associations

Service-Id : 1570 (VPLS) Customer-Id : 1
- SAP : lag-4:1570

SAP Egress Policy Id : 16401

Associations

Service-Id : 1640 (VPLS) Customer-Id : 1
- SAP : lag-4:1640

SAP Egress Policy Id : 20001

Network

Network Policy Id : 50

Interface Association

Interface	: ip-192.162.105.4		
IP Addr.	: 192.162.105.4/24	Port Id	: 1/1/23
Interface	: ip-192.162.20.4		
IP Addr.	: 192.162.20.4/24	Port Id	: lag-3
Interface	: ip-192.162.45.4		
IP Addr.	: 192.162.45.4/24	Port Id	: lag-5
Interface	: ip-192.162.80.4		
IP Addr.	: 192.162.80.4/24	Port Id	: 1/1/22
Network Policy Id	: 550		

Interface Association

Interface	: ip-192.162.100.4		
IP Addr.	: 192.162.100.4/24	Port Id	: 1/1/23
Interface	: ip-192.162.40.4		
IP Addr.	: 192.162.40.4/24	Port Id	: lag-5
Interface	: ip-192.162.65.4		
IP Addr.	: 192.162.65.4/24	Port Id	: 1/1/13
Network Policy Id	: 1050		

Interface Association

No Interface Association Found.

Network Policy Id : 1550

```

-----
Interface Association
-----
No Interface Association Found.

Network Policy Id           : 2050

-----
Interface Association
-----
No Interface Association Found.

=====

*A:SAS-X-C>config>qos# show qos remark-policy 500 detail

=====
QoS Remarking Policies
=====
-----
Remark Policy-id   : 500           Type           : dot1p-lsp-exp-shared
Description        : (Not Specified)

-----
FC Name      dot1P / LSP EXP      dot1P / LSP EXP
              In Value             Out Value
-----
be           3                    6
l2           1                    4
af           6                    1
l1           7                    2
h2           0                    3
ef           5                    0
h1           4                    7
nc           2                    5

-----
Associations
-----
SAP Egress
-----
SAP Egress Policy Id       : 5001

-----
Associations
-----
Service-Id   : 500 (VPLS)           Customer-Id : 1
- SAP : lag-2:500

SAP Egress Policy Id       : 5701

-----
Associations
-----
Service-Id   : 570 (VPLS)           Customer-Id : 1
- SAP : lag-2:570

SAP Egress Policy Id       : 6401

```

```

-----
Associations
-----
Service-Id      : 640 (VPLS)                Customer-Id   : 1
- SAP : lag-2:640

SAP Egress Policy Id      : 10001

-----
Associations
-----
Service-Id      : 1000 (VPLS)                Customer-Id   : 1
- SAP : lag-2:1000

SAP Egress Policy Id      : 10701

-----
Associations
-----
Service-Id      : 1070 (VPLS)                Customer-Id   : 1
- SAP : lag-2:1070

SAP Egress Policy Id      : 11401

-----
Associations
-----
Service-Id      : 1140 (VPLS)                Customer-Id   : 1
- SAP : lag-2:1140

SAP Egress Policy Id      : 15001

-----
Associations
-----
Service-Id      : 1500 (VPLS)                Customer-Id   : 1
- SAP : lag-4:1500

SAP Egress Policy Id      : 15701

-----
Associations
-----
Service-Id      : 1570 (VPLS)                Customer-Id   : 1
- SAP : lag-4:1570

SAP Egress Policy Id      : 16401

-----
Associations
-----
Service-Id      : 1640 (VPLS)                Customer-Id   : 1
- SAP : lag-4:1640

SAP Egress Policy Id      : 20001

-----
Associations
-----

```

```

-----
Service-Id      : 2000 (Epipe)                Customer-Id   : 1
- SAP : 1/1/2:2000

SAP Egress Policy Id      : 20701

-----
Network
-----
Network Policy Id      : 50

-----
Interface Association
-----
Interface      : ip-192.162.105.4
IP Addr.       : 192.162.105.4/24              Port Id      : 1/1/23
Interface      : ip-192.162.20.4
IP Addr.       : 192.162.20.4/24              Port Id      : lag-3
Interface      : ip-192.162.45.4
IP Addr.       : 192.162.45.4/24              Port Id      : lag-5
Interface      : ip-192.162.80.4
IP Addr.       : 192.162.80.4/24              Port Id      : 1/1/22
Network Policy Id      : 550

-----
Interface Association
-----
Interface      : ip-192.162.100.4
IP Addr.       : 192.162.100.4/24              Port Id      : 1/1/23
Interface      : ip-192.162.40.4
IP Addr.       : 192.162.40.4/24              Port Id      : lag-5
Interface      : ip-192.162.65.4
IP Addr.       : 192.162.65.4/24              Port Id      : 1/1/13
Network Policy Id      : 1050

-----
Interface Association
-----
No Interface Association Found.

Network Policy Id      : 1550

-----
Interface Association
-----
No Interface Association Found.

Network Policy Id      : 2050

-----
Interface Association
-----
No Interface Association Found.

=====
*A:SAS-X-C>config>qos# show qos remark-policy 500 detail

```

```

=====
QoS Remarking Policies
=====
-----
Remark Policy-id      : 500          Type          : dot1p-lsp-exp-shared
Description           : (Not Specified)
-----

FC Name      dot1P / LSP EXP      dot1P / LSP EXP
              In Value            Out Value
-----
be           3                    6
l2           1                    4
af           6                    1
l1           7                    2
h2           0                    3
ef           5                    0
h1           4                    7
nc           2                    5
-----

Associations
-----
SAP Egress
-----
SAP Egress Policy Id      : 5001
-----

Associations
-----
Service-Id      : 500 (VPLS)          Customer-Id : 1
- SAP : lag-2:500
SAP Egress Policy Id      : 5701
-----

Associations
-----
Service-Id      : 570 (VPLS)          Customer-Id : 1
- SAP : lag-2:570
SAP Egress Policy Id      : 6401
-----

Associations
-----
Service-Id      : 640 (VPLS)          Customer-Id : 1
- SAP : lag-2:640
SAP Egress Policy Id      : 10001
-----

Network
-----
Network Policy Id      : 50
-----

```

Interface Association

```
-----  
Interface      : ip-192.162.105.4  
IP Addr.       : 192.162.105.4/24          Port Id       : 1/1/23  
Interface      : ip-192.162.20.4  
IP Addr.       : 192.162.20.4/24          Port Id       : lag-3  
Interface      : ip-192.162.45.4  
IP Addr.       : 192.162.45.4/24          Port Id       : lag-5  
Interface      : ip-192.162.80.4  
IP Addr.       : 192.162.80.4/24          Port Id       : 1/1/22  
Network Policy Id      : 550  
-----
```

Interface Association

```
-----  
Interface      : ip-192.162.100.4  
IP Addr.       : 192.162.100.4/24          Port Id       : 1/1/23  
Interface      : ip-192.162.40.4  
IP Addr.       : 192.162.40.4/24          Port Id       : lag-5  
Interface      : ip-192.162.65.4  
IP Addr.       : 192.162.65.4/24          Port Id       : 1/1/13  
Network Policy Id      : 1050  
-----
```

Interface Association

No Interface Association Found.

Network Policy Id : 1550

Interface Association

No Interface Association Found.

Network Policy Id : 2050

Interface Association

No Interface Association Found.

```
-----  
=====
```

Multipoint Bandwidth Management

In This Section

This section provides information on multipoint bandwidth management using the command line interface.

Topics in this section include:

- [Overview on page 344](#)
- [Configuration Guidelines on page 345](#)

Overview

In the 7210 SAS, multicast, unknown unicast, and broadcast traffic types are en-queued into a set of eight ingress queues. SAP ingress traffic and network traffic which are been policed by multipoint SAP ingress meters are en-queued into these queues before being replicated to appropriate egress port. There are a set of eight queues per node which are shared by all the services. Packets are en-queued to the appropriate queue based on the Forwarding Class (FC) assigned to the packet by the SAP ingress classification.

The Multipoint Bandwidth Management CLI commands are used to configure a bandwidth policy to manage different traffic types like broadcast, unknown unicast and multicast traffic. The aggregate ingress rate can be configured to control the amount of multipoint traffic per node. For each of the queues the user can specify the CIR/PIR and CBS/MBS to control the amount of traffic and packet buffers allocated per FC respectively.

If a policy is not configured explicitly by the user, a default multipoint ingress policy is used by the system. A default slope-policy is used by the system, but this policy cannot be configured by the user.

The default slope policy for a multipoint queue:

```
high-slope
  start-avg 70
  max-avg 90
  max-prob 75
no shutdown
exit
low-slope
  start-avg 50
  max-avg 75
  max-prob 75
no shutdown
exit
```

Configuration Guidelines

- The FC to multipoint queue map cannot be configured by the user. It is defined by the system.
- The maximum total amount of ingress multicast traffic which can be replicated is approximately 10Gbps.
- The maximum total amount of egress multicast traffic (after replication) is limited to approximately 23Gbps.
- The ingress multipoint queues are scheduled in strict priority order, with FC NC being scheduled first and FC BE being scheduled last.
- WRED slopes for the queues cannot be configured by the users.

Multipoint Bandwidth Management Command Reference

Command Hierarchies

Configuration Commands

```
config
— multipoint-management
  — [no] bandwidth-policy <policy-name>
    — [no] description description-string
    — [no] ingress-aggregate-rate <megabits-per-second>
    — queue <queue-id>
      — [no] adaptation-rule [pir <adaptation-rule>] [cir <adaptation-rule>]
      — [no] rate cir <cir-percent>[ pir <pir-percent>]
      — [no] cbs <in kilo bytes>
      — [no] mbs <in kilo bytes>

config
— system
  — multipoint-management
    — [no] bandwidth-policy <policy-name>
```

SHOW COMMANDS

```
show
— system
  — multipoint-management <detail>
— multipoint-management
  — bandwidth-policy
  — bandwidth-policy <detail>
  — bandwidth-policy <policy-name>
  — bandwidth-policy <policy-name> <detail>
```

CLEAR COMMANDS

```
clear
— system
— multipoint-management statistics
```

Configuration Commands

Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>system>multipoint-management>bandwidth-policy
Description	This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the context in the configuration file. The no form of this command removes any description string from the context.
Default	No description is associated with the configuration context.
Parameters	<i>description-string</i> — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Multipoint Bandwidth Management Commands

multipoint-management

Syntax	multipoint-management
Context	config
Description	This command enables the context to configure Multipoint Bandwidth Management.

bandwidth-policy

Syntax	[no] bandwidth-policy <i><policy-name></i>
Context	config>multipoint-management config>system>multipoint-management
Description	This command specifies the name of the multipoint bandwidth management policy.
Parameters	<i>policy name</i> — The name of the multipoint bandwidth management policy. Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces and so on), the entire string must be enclosed within double quotes.

ingress-aggregate-rate

Syntax	[no] ingress-aggregate-rate <i><megabits-per-second></i>
Context	config>multipoint-management>bandwidth-policy
Description	This command configures the total multipoint ingress traffic rate.
Default	10000
Parameters	<i>megabits-per-second</i> — Specifies the ingress aggregate rate in Mbps. Values 1 — 10000

queue

Syntax	queue <queue-id>
Context	config>multipoint-management>bandwidth-policy
Description	This command provides the context to configure the multipoint queue parameters. Eight queues are available for the user to configure. The forwarding classes (FCs) assignment to queues is given in Table 38 .
Default	None
Parameters	queue id — Specifies the queue ID. Values 1- 8

Table 38: FC Queue Table

FC	Queue
BE	Queue1
L2	Queue2
AF	Queue3
L1	Queue4
H2	Queue5
EF	Queue6
H1	Queue7
NC	Queue8

adaptation-rule

Syntax	[no] adaptation-rule [cir <adaptation-rule>] [pir <adaptation-rule>]
Context	config>multipoint-management>bandwidth-policy>queue
Description	This command defines the method used by the system to derive the operational CIR and PIR settings when the queue’s rate is provisioned in hardware. For the CIR and PIR parameters, individually the system attempts to find the best operational rate depending on the defined constraint and the applicable hardware constraints. The no form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific adaptation-rule is removed, the default constraints for cir and pir apply.

Default closest

Parameters *adaptation-rule* — Specifies the adaptation rule to be used while computing the operational CIR or PIR value.

max — The **max** (maximum) option is mutually exclusive with the min and closest options.

min — The **min** (minimum) option is mutually exclusive with the max and closest options.

closest — The **closest** parameter is mutually exclusive with the min and max parameter.

rate

Syntax **[no] rate cir <cir-percent> [pir <pir-percent>]**

Context config>multipoint-management>bandwidth-policy>queue

Description This command defines the administrative PIR and CIR parameters for the queue. The CIR and PIR rates are specified in percentage. The system determines the operational rate as a percentage of the ingress-aggregate-rate configured.

The **no** form of the command restores the CIR and PIR rates to default values.

Default 100

Parameters *cir-percent* — Defines the percentage of the guaranteed rate allowed for the queue. When the rate command is executed, a valid CIR setting must be explicitly defined. When the rate command has not been executed, the default CIR of 0 is assumed. Fractional values are not allowed and must be given as a positive integer.

The actual CIR rate is dependent on the queue's adaptation-rule parameters and the actual hardware where the queue is provisioned.

Values 0 — 100

pir-percent — Defines the percentage of the maximum rate allowed for the queue. When the rate command is executed, the PIR setting is optional. When the rate command has not been executed, or the PIR parameter is not explicitly specified, the default PIR of 100 is assumed. Fractional values are not allowed and must be given as a positive integer.

Values 1 — 100 percent

cbs

Syntax **[no] cbs <in kilo bytes>**

Context config>multipoint-management>bandwidth-policy>queue

Description The Committed Burst Size (CBS) specifies the amount of reserved buffers for a specified queue. The value is given in kilo-bytes.

The CBS for a queue is used to determine whether the queue has exhausted its reserved buffers while en-queuing packets. Once the queue has exceeded the amount of buffers considered in reserve it must contend with other queues for the available shared buffer space within the system buffer pool. Access to the shared pool is controlled through Random Early Detection (RED) application. Oversubscription of CBS is not allowed.

Two RED slopes are maintained for each queue. A high priority slope is used by in-profile packets. A low priority slope is used by out-of-profile packets. The RED slopes are not user configurable.

The no form of this command returns the CBS size for the queue to the default value.

Default Queue1— 1024KB
Queue[2 to 6] — 2048 KB
Queue[7 to 8] —512KB

Parameters <in kilo bytes> — pecifies the CBS values in kilo bytes.
Values 0—131072

mbs

Syntax [no] mbs <in kilo bytes>

Context config>multipoint-management>bandwidth-policy>queue

Description The Maximum Burst Size (MBS) command is used to specify maximum amount of buffers that can be used by a particular queue from the shared pool of buffers. This value is specified in kilobytes.

The MBS value is used by a queue to determine whether it has exhausted its total allowed buffers while en-queuing packets. Once the queue has exceeded its maximum amount of buffers, all packets are discarded until the queue transmits a packet. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet.

The **no** form of the command returns the MBS size for the queue to the default value.

Default Queue[1 to 6] — 6144 KB
Queue[7 to 8] — 2048KB

Parameters <in kilo bytes> — Specifies the MBS values in kilo bytes.
Values 0—131072

Show Commands

multipoint-management

- Syntax** **multipoint-management**
- Context** show>system
- Description** This command displays the multipoint management information.

bandwidth-policy

- Syntax** **bandwidth-policy**
bandwidth-policy <detail>
bandwidth-policy <policy-name>
bandwidth-policy <policy-name> <detail>
- Context** show>multipoint-management
- Description** This command displays the multipoint management information for the bandwidth policies.
- Parameters** *policy-name* — The name of the policy.
detail — Displays detailed information of the multipoint-management policy.

Output **Multipoint-management Policy Output Fields** — The following table describes Multipoint-management policy output fields.

Table 39: Show Multipoint-management Policy Output Fields

Label	Description
Policy	Displays the policy name of the Multipoint-management policy.
Description	A string that identifies the policy’s context in the configuration file.
Ingr. Aggr. Rate	Displays the Ingress aggregate rate.
Queue	Displays the queue ID of the queue.
CBS	Displays the configured CBS value.
CIR	Displays the committed information rate.
PIR	Displays the peak information rate.
Mgmt Plcy	Displays the configured management policy.

Table 39: Show Multipoint-management Policy Output Fields (Continued)

Label	Description
MBS	Displays the configured MBS value.
CIR Adaptation Rule	Displays the adaptation rule in use.
PIR Adaptation Rule	Displays the adaptation rule in use.
Queue Statistics	
Egress Queue	Displays the egress queue ID and the associated forwarding class.
Fwd Stats	Displays the forwarding statistics in octets and packets.
Drop Stats	Displays the drop statistics in octets and packets.

Sample Output

```
*A:Dut-G>show>mpoint-mgmt># bandwidth-policy detail
```

```
=====
Bandwidth Policy Details
=====
-----
Policy          : abc
Description     :                               Ingr. Aggr. Rate: Default
-----
Queue 1         :                               Mgmt Plcy       : None
CBS              : default                       MBS                 : default
CIR              : default                       CIR Adaptation Rule: closest
PIR              : default                       PIR Adaptation Rule: closest
Queue 2         :                               Mgmt Plcy       : None
CBS              : default                       MBS                 : default
CIR              : default                       CIR Adaptation Rule: closest
PIR              : default                       PIR Adaptation Rule: closest
Queue 3         :                               Mgmt Plcy       : None
CBS              : default                       MBS                 : default
CIR              : default                       CIR Adaptation Rule: closest
PIR              : default                       PIR Adaptation Rule: closest
Queue 4         :                               Mgmt Plcy       : None
CBS              : default                       MBS                 : default
CIR              : default                       CIR Adaptation Rule: closest
PIR              : default                       PIR Adaptation Rule: closest
Queue 5         :                               Mgmt Plcy       : None
CBS              : default                       MBS                 : default
CIR              : default                       CIR Adaptation Rule: closest
PIR              : default                       PIR Adaptation Rule: closest
Queue 6         :                               Mgmt Plcy       : None
CBS              : default                       MBS                 : default
CIR              : default                       CIR Adaptation Rule: closest
PIR              : default                       PIR Adaptation Rule: closest
```

```

Queue 7      :                               Mgmt Plcy      : None
CBS          : default                       MBS             : default
CIR          : default                       CIR Adaptation Rule: closest
PIR          : default                       PIR Adaptation Rule: closest
Queue 8      :                               Mgmt Plcy      : None
CBS          : default                       MBS             : default
CIR          : default                       CIR Adaptation Rule: closest
PIR          : default                       PIR Adaptation Rule: closest
-----
Policy       : default
Description   :                               Ingr. Aggr. Rate: Default
-----
Queue 1      :                               Mgmt Plcy      : None
CBS          : default                       MBS             : default
CIR          : default                       CIR Adaptation Rule: closest
PIR          : default                       PIR Adaptation Rule: closest
Queue 2      :                               Mgmt Plcy      : None
CBS          : default                       MBS             : default
CIR          : default                       CIR Adaptation Rule: closest
PIR          : default                       PIR Adaptation Rule: closest
Queue 3      :                               Mgmt Plcy      : None
CBS          : default                       MBS             : default
CIR          : default                       CIR Adaptation Rule: closest
PIR          : default                       PIR Adaptation Rule: closest
Queue 4      :                               Mgmt Plcy      : None
CBS          : default                       MBS             : default
CIR          : default                       CIR Adaptation Rule: closest
PIR          : default                       PIR Adaptation Rule: closest
Queue 5      :                               Mgmt Plcy      : None
CBS          : default                       MBS             : default
CIR          : default                       CIR Adaptation Rule: closest
PIR          : default                       PIR Adaptation Rule: closest
Queue 6      :                               Mgmt Plcy      : None
CBS          : default                       MBS             : default
CIR          : default                       CIR Adaptation Rule: closest
PIR          : default                       PIR Adaptation Rule: closest
Queue 7      :                               Mgmt Plcy      : None
CBS          : default                       MBS             : default
CIR          : default                       CIR Adaptation Rule: closest
PIR          : default                       PIR Adaptation Rule: closest
Queue 8      :                               Mgmt Plcy      : None
CBS          : default                       MBS             : default
CIR          : default                       CIR Adaptation Rule: closest
PIR          : default                       PIR Adaptation Rule: closest
=====
Bandwidth Policies : 2
=====

```

*A:Dut-G>show# multipoint-management bandwidth-policy abc

```

=====
Bandwidth Policies
=====
Bw Policy      Description                               Ingr. Aggr. *
-----
abc            Default
=====

```

```

Bandwidth Policies : 1
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-G>show#

*A:Dut-G>show>mpoint-mgmt># bandwidth-policy abc detail

=====
Bandwidth Policy Details
=====
-----
Policy          : abc
Description     :                               Ingr. Aggr. Rate: Default
-----
Queue 1         :                               Mgmt Plcy       : None
CBS              : default                       MBS                 : default
CIR              : default                       CIR Adaptation Rule: closest
PIR              : default                       PIR Adaptation Rule: closest
Queue 2         :                               Mgmt Plcy       : None
CBS              : default                       MBS                 : default
CIR              : default                       CIR Adaptation Rule: closest
PIR              : default                       PIR Adaptation Rule: closest
Queue 3         :                               Mgmt Plcy       : None
CBS              : default                       MBS                 : default
CIR              : default                       CIR Adaptation Rule: closest
PIR              : default                       PIR Adaptation Rule: closest
Queue 4         :                               Mgmt Plcy       : None
CBS              : default                       MBS                 : default
CIR              : default                       CIR Adaptation Rule: closest
PIR              : default                       PIR Adaptation Rule: closest
Queue 5         :                               Mgmt Plcy       : None
CBS              : default                       MBS                 : default
CIR              : default                       CIR Adaptation Rule: closest
PIR              : default                       PIR Adaptation Rule: closest
Queue 6         :                               Mgmt Plcy       : None
CBS              : default                       MBS                 : default
CIR              : default                       CIR Adaptation Rule: closest
PIR              : default                       PIR Adaptation Rule: closest
Queue 7         :                               Mgmt Plcy       : None
CBS              : default                       MBS                 : default
CIR              : default                       CIR Adaptation Rule: closest
PIR              : default                       PIR Adaptation Rule: closest
Queue 8         :                               Mgmt Plcy       : None
CBS              : default                       MBS                 : default
CIR              : default                       CIR Adaptation Rule: closest
PIR              : default                       PIR Adaptation Rule: closest
=====
Bandwidth Policies : 1
=====
*A:Dut-G>show>mpoint-mgmt>#

*A:Dut-G>show>system# multipoint-management detail

=====
System Multipoint Bandwidth Policy Details
=====
-----
Policy          : abc

```

Description	:	Ingr. Aggr. Rate: Default
Queue 1	:	Mgmt Plcy : None
CBS	:	MBS : default
CIR	:	CIR Adaptation Rule: closest
PIR	:	PIR Adaptation Rule: closest
Queue 2	:	Mgmt Plcy : None
CBS	:	MBS : default
CIR	:	CIR Adaptation Rule: closest
PIR	:	PIR Adaptation Rule: closest
Queue 3	:	Mgmt Plcy : None
CBS	:	MBS : default
CIR	:	CIR Adaptation Rule: closest
PIR	:	PIR Adaptation Rule: closest
Queue 4	:	Mgmt Plcy : None
CBS	:	MBS : default
CIR	:	CIR Adaptation Rule: closest
PIR	:	PIR Adaptation Rule: closest
Queue 5	:	Mgmt Plcy : None
CBS	:	MBS : default
CIR	:	CIR Adaptation Rule: closest
PIR	:	PIR Adaptation Rule: closest
Queue 6	:	Mgmt Plcy : None
CBS	:	MBS : default
CIR	:	CIR Adaptation Rule: closest
PIR	:	PIR Adaptation Rule: closest
Queue 7	:	Mgmt Plcy : None
CBS	:	MBS : default
CIR	:	CIR Adaptation Rule: closest
PIR	:	PIR Adaptation Rule: closest
Queue 8	:	Mgmt Plcy : None
CBS	:	MBS : default
CIR	:	CIR Adaptation Rule: closest
PIR	:	PIR Adaptation Rule: closest

Queue Statistics

	Packets	Octets
Egress Queue 1 (be)		
Fwd Stats :	0	0
Drop Stats :	0	0
Egress Queue 2 (l2)		
Fwd Stats :	0	0
Drop Stats :	0	0
Egress Queue 3 (af)		
Fwd Stats :	0	0
Drop Stats :	0	0
Egress Queue 4 (l1)		
Fwd Stats :	0	0
Drop Stats :	0	0
Egress Queue 5 (h2)		
Fwd Stats :	0	0
Drop Stats :	0	0
Egress Queue 6 (ef)		

Fwd Stats	:	79	7426
Drop Stats	:	0	0
Egress Queue 7 (h1)			
Fwd Stats	:	0	0
Drop Stats	:	0	0
Egress Queue 8 (nc)			
Fwd Stats	:	109036873005	8182107400620
Drop Stats	:	4283683389	351310933698

=====

*A:Dut-G>show>system#

Clear Commands

multipoint-management

Syntax **multipoint-management** *statistics*

Context clear>system

Description This command clears the queue counters.

Standards and Protocol Support

Standards Compliance

IEEE 802.1ab-REV/D3 Station and Media Access Control Connectivity Discovery
IEEE 802.1D Bridging
IEEE 802.1p/Q VLAN Tagging
IEEE 802.1s Multiple Spanning Tree
IEEE 802.1w Rapid Spanning Tree Protocol
IEEE 802.1X Port Based Network Access Control
IEEE 802.1ad Provider Bridges
IEEE 802.1ah Provider Backbone Bridges
IEEE 802.1ag Service Layer OAM
IEEE 802.3ah Ethernet in the First Mile
IEEE 802.3 10BaseT
IEEE 802.3ad Link Aggregation
IEEE 802.3ae 10Gbps Ethernet
IEEE 802.3ah Ethernet OAM
IEEE 802.3u 100BaseTX
IEEE 802.3z 1000BaseSX/LX ITU-T Y.1731 OAM functions and mechanisms for Ethernet based networks draft-ietf-disman-alarm-mib-04.txt IANA-IFType-MIB
IEEE8023-LAG-MIB ITU-T G.8032 Ethernet Ring Protection Switching (version 2)

Protocol Support

BGP

RFC 1397 BGP Default Route Advertisement
RFC 1772 Application of BGP in the Internet
RFC 1997 BGP Communities Attribute
RFC 2385 Protection of BGP Sessions via MD5
RFC 2439 BGP Route Flap Dampening
RFC 2547 bis BGP/MPLS VPNs draft-ietf-idr-rfc2858bis-09.txt.
RFC 2918 Route Refresh Capability for BGP-4
RFC 3392 Capabilities Advertisement with BGP4
RFC 4271 BGP-4 (previously RFC 1771)

RFC 4360 BGP Extended Communities Attribute
RFC 4364 BGP/MPLS IP Virtual Private Networks (VPNs)(previously RFC 2547bis BGP/MPLS VPNs)
RFC 4456 BGP Route Reflection:Alternative to Full-mesh IBGP (Previously RFC 1966 & 2796)
RFC 4760 Multi-protocol Extensions for BGP
RFC 4893 BGP Support for Four-octet AS Number Space

CIRCUIT EMULATION

RFC 4553 Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)
RFC 5086 Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)
RFC 5287 Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks

DHCP

RFC 2131 Dynamic Host Configuration Protocol (REV)

DIFFERENTIATED SERVICES

RFC 2474 Definition of the DS Field the IPv4 and IPv6 Headers (Rev)
RFC 2597 Assured Forwarding PHB Group (rev3260)
RFC 2598 An Expedited Forwarding PHB
RFC 2697 A Single Rate Three Color Marker
RFC 2698 A Two Rate Three Color Marker
RFC 4115 A Differentiated Service Two-Rate, Three-Color Marker with Efficient Handling of in-Profile Traffic

IPv6

RFC 2460 Internet Protocol, Version 6 (IPv6) Specification
RFC 2461 Neighbor Discovery for IPv6

RFC 2463 Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 Specification
RFC 2464 Transmission of IPv6 Packets over Ethernet Networks

IS-IS

RFC 1142 OSI IS-IS Intra-domain Routing Protocol (ISO 10589)
RFC 1195 Use of OSI IS-IS for routing in TCP/IP & dual environments
RFC 2763 Dynamic Hostname Exchange for IS-IS
RFC 2966 Domain-wide Prefix Distribution with Two-Level IS-IS
RFC 2973 IS-IS Mesh Groups
RFC 3373 Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies
RFC 3567 Intermediate System to Intermediate System (ISIS) Cryptographic Authentication
RFC 3719 Recommendations for Interoperable Networks using IS-IS
RFC 3784 Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)
RFC 3787 Recommendations for Interoperable IP Networks
RFC 3847 Restart Signaling for IS-IS – GR helper

LDP

RFC 3036 LDP Specification
RFC 3037 LDP Applicability
RFC 3478 Graceful Restart Mechanism for LDP — GR helper
RFC 5283 LDP extension for Inter-Area LSP draft-jork-ldp-igp-sync-03.txt

MPLS

RFC 3031 MPLS Architecture
RFC 3032 MPLS Label Stack Encoding (REV3443))
RFC 4379 Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures

RFC 4182 Removing a Restriction on the use of MPLS Explicit NULL
draft-ietf-mpls-lsr-mib-06.txt
draft-ietf-mpls-te-mib-04.txt
draft-ietf-mpls-ldp-mib-07.txt

Multicast

RFC 1112 Host Extensions for IP Multicasting (Snooping)
RFC 2236 Internet Group Management Protocol, (Snooping)
RFC 3376 Internet Group Management Protocol, Version 3 (Snooping) [Only in 7210 SAS-M access-uplink mode]

NETWORK MANAGEMENT

ITU-T X.721: Information technology- OSI-Structure of Management Information
ITU-T X.734: Information technology- OSI-Systems Management: Event Report Management Function
M.3100/3120 Equipment and Connection Models
TMF 509/613 Network Connectivity Model
RFC 1157 SNMPv1
RFC 1215 A Convention for Defining Traps for use with the SNMP
RFC 1907 SNMPv2-MIB
RFC 2011 IP-MIB
RFC 2012 TCP-MIB
RFC 2013 UDP-MIB
RFC 2096 IP-FORWARD-MIB
RFC 2138 RADIUS
RFC 2206 RSVP-MIB
RFC 2571 SNMP-FRAMEWORKMIB
RFC 2572 SNMP-MPD-MIB
RFC 2573 SNMP-TARGET-&-NOTIFICATION-MIB
RFC 2574 SNMP-USER-BASEDSMMIB
RFC 2575 SNMP-VIEW-BASEDACM-MIB
RFC 2576 SNMP-COMMUNITY-MIB
RFC 2665 EtherLike-MIB
RFC 2819 RMON-MIB
RFC 2863 IF-MIB
RFC 2864 INVERTED-STACK-MIB
RFC 3014 NOTIFICATION-LOGMIB
RFC 3164 Syslog

RFC 3273 HCRMON-MI
RFC 3411 An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks
RFC 3412 - Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
RFC 3413 - Simple Network Management Protocol (SNMP) Applications
RFC 3414 - User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)
RFC 3418 - SNMP MIB

OSPF

RFC 1765 OSPF Database Overflow
RFC 2328 OSPF Version 2
RFC 2370 Opaque LSA Support
RFC 3101 OSPF NSSA Option
RFC 3137 OSPF Stub Router Advertisement
RFC 3623 Graceful OSPF Restart – GR helper
RFC 3630 Traffic Engineering (TE) Extensions to OSPF Version 2

RSVP-TE

RFC 2430 A Provider Architecture DiffServ & TE
RFC 2702 Requirements for Traffic Engineering over MPLS
RFC2747 RSVP Cryptographic Authentication
RFC3097 RSVP Cryptographic Authentication
RFC 3209 Extensions to RSVP for Tunnels
RFC 4090 Fast reroute Extensions to RSVP-TE for LSP Tunnels draft-ietf-ccamp-mpls-graceful-shutdown-06 Graceful Shutdown in GMPLS Traffic Engineering Networks

PSEUDO-WIRE

RFC 3985 Pseudo Wire Emulation Edge-to-Edge (PWE3)
RFC 4385 Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN

RFC 3916 Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)
RFC 4448 Encapsulation Methods for Transport of Ethernet over MPLS Networks (draft-ietf-pwe3-ethernet-encap-11.txt)
RFC 4446 IANA Allocations for PWE3
RFC 4447 Pseudowire Setup and Maintenance Using LDP (draft-ietf-pwe3-control-protocol-17.txt)
RFC 5085, Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires
draft-ietf-l2vpn-vpws-iw-oam-02.txt
draft-ietf-pwe3-oam-msg-map-05.txt
draft-ietf-pwe3-ms-pw-arch-02.txt
draft-ietf-pwe3-segmented-pw-05.txt
draft-hart-pwe3-segmented-pw-vcv-02.txt
draft-muley-dutta-pwe3-redundancy-bit-02.txt
draft-muley-pwe3-redundancy-02.txt

RADIUS

RFC 2865 Remote Authentication Dial In User Service
RFC 2866 RADIUS Accounting

SSH

draft-ietf-secsh-architecture.txt SSH Protocol Architecture
draft-ietf-secsh-userauth.txt SSH Authentication Protocol
draft-ietf-secsh-transport.txt SSH Transport Layer Protocol
draft-ietf-secsh-connection.txt SSH Connection Protocol
draft-ietf-secsh- newmodes.txt SSH Transport Layer Encryption Modes

TACACS+

draft-grant-tacacs-02.txt

TCP/IP

RFC 768 UDP
RFC 1350 The TFTP Protocol
RFC 791 IP
RFC 792 ICMP
RFC 793 TCP
RFC 826 ARP
RFC 854 Telnet
RFC 1519 CIDR

RFC 1812 Requirements for IPv4 Routers
 RFC 2347 TFTP option Extension
 RFC 2328 TFTP Blocksize Option
 RFC 2349 TFTP Timeout Interval and Transfer Size option

Timing

ITU-T G.781 Telecommunication Standardization Section of ITU, Synchronization layer functions, issued 09/2008
 ITU-T G.813 Telecommunication Standardization Section of ITU, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003.
 GR-1244-CORE Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005
 ITU-T G.8261 Telecommunication Standardization Section of ITU, Timing and synchronization aspects in packet networks, issued 04/2008.
 ITU-T G.8262 Telecommunication Standardization Section of ITU, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007.
 ITU-T G.8264 Telecommunication Standardization Section of ITU, Distribution of timing information through packet networks, issued 10/2008.

VPLS

RFC 4762 Virtual Private LAN Services Using LDP (previously draft-ietf-l2vpn-vpls-ldp-08.txt)

Proprietary MIBs

ALCATEL-IGMP-SNOOPING-MIB.mib
 TIMETRA-CAPABILITY-7210-SAS-M-V1v0.mib
 TIMETRA-CHASSIS-MIB.mib
 TIMETRA-CLEAR-MIB.mib
 TIMETRA-DOT3-OAM-MIB.mib
 TIMETRA-FILTER-MIB.mib
 TIMETRA-GLOBAL-MIB.mib
 TIMETRA-IEEE8021-CFM-MIB.mib
 TIMETRA-LAG-MIB.mib
 TIMETRA-LOG-MIB.mib

TIMETRA-MIRROR-MIB.mib
 TIMETRA-NTP-MIB.mib
 TIMETRA-OAM-TEST-MIB.mib
 TIMETRA-PORT-MIB.mib
 TIMETRA-QOS-MIB.mib
 TIMETRA-SAS-ALARM-INPUT-MIB.mib
 TIMETRA-SAS-IEEE8021-CFM-MIB.mib
 TIMETRA-SAS-GLOBAL-MIB.mib
 TIMETRA-SAS-PORT-MIB.mib
 TIMETRA-SAS-QOS-MIB.mib
 TIMETRA-SAS-SYSTEM-MIB.mib
 TIMETRA-SAS-SERV-MIB.mib
 TIMETRA-SAS-VRTR-MIB.mib
 TIMETRA-SCHEDULER-MIB.mib
 TIMETRA-SECURITY-MIB.mib
 TIMETRA-SERV-MIB.mib
 TIMETRA-SYSTEM-MIB.mib
 TIMETRA-TC-MIB.mib
 TIMETRA-ISIS-MIB.mib
 TIMETRA-ROUTE-POLICY-MIB.mib
 TIMETRA-MPLS-MIB.mib
 TIMETRA-RSVP-MIB.mib
 TIMETRA-LDP-MIB.mib
 TIMETRA-VRTR-MIB.mib

INDEX

Q

QoS

- overview 18
 - policies 19
 - policy entities 64
- access egress
 - overview 262
 - configuring
 - access egress policies 263
 - applying policies 264
 - command reference 269, 297
 - default values 265
- frame-based accounting
 - overview 78
 - configuring
 - disable 80
 - enable 80
- network policies
 - overview 23
 - configuring
 - basic 96
 - command reference 115
 - default policy values 99
 - overview 88
- network queue policies
 - overview 28
 - adaptation rule 38
 - CBS 40
 - CIR 36
 - PIR 37
 - queue ID 35
 - configuring
 - applying to network ingress port 154
 - basic 152
 - default policy values 155
 - overview 150
 - command reference 163
- SAP policies
 - overview
 - egress policies 46
 - ingress policies 41
 - configuring

- applying to services 220
- basic 186
- command reference 225
- ingress policy 187
- overview 176, 278
- slope policies
 - overview 53, 304, 320, 344
 - RED slopes 54
 - shared buffer utilization 56
- configuring
 - basic 305, 322

